arXiv:1408.6908v2 [cs.AI] 10 Oct 2014

AI Evaluation: past, present and future*

José Hernández-Orallo
DSIC, Universitat Politècnica de València, Spain
jorallo@dsic.upv.es

October 13, 2014

Abstract

Artificial intelligence develops techniques and systems whose performance must be evaluated on a regular basis in order to certify and foster progress in the discipline. We will describe and critically assess the different ways AI systems are evaluated. We first focus on the traditional *task*-oriented evaluation approach. We see that black-box (behavioural evaluation) is becoming more and more common, as AI systems are becoming more complex and unpredictable. We identify three kinds of evaluation: human discrimination, problem benchmarks and peer confrontation. We describe the limitations of the many evaluation settings and competitions in these three categories and propose several ideas for a more systematic and robust evaluation. We then focus on a less customary (and challenging) *ability*-oriented evaluation approach, where a system is characterised by its (cognitive) abilities, rather than by the tasks it is designed to solve. We discuss several possibilities: the adaptation of cognitive tests used for humans and animals, the development of tests derived from algorithmic information theory or more general approaches under the perspective of universal psychometrics.

Keywords: AI evaluation, AI competitions, benchmark evaluation, sampling, narrow vs general AI, measurement, universal psychometrics, Turing Test.

1 Introduction

The evaluation of any discipline must necessarily be linked to the purpose of the discipline. What is the purpose of artificial intelligence (AI)? McCarthy's pristine definition of AI sets this unambiguously: "[AI is] the science and engineering of making intelligent machines" [123]. As a consequence, AI evaluation should focus on evaluating the intelligence of the artefacts it builds. However, as we will further discuss below, 'intelligence tests' (of whatever kind) are not the everyday evaluation approach for AI. The explanation for this is that most AI research is better identified by Minsky's more pragmatic definition: "[AI is] the science of making machines capable of performing tasks that would require intelligence if done by [humans]" [129, p.v]. As a result, AI evaluation focusses on checking whether machines do these *tasks* well.

This has led to an important anomaly of AI. AI artefacts solve these tasks *without featuring intelligence*. Paradoxically, this is one of the reasons of AI success. Systems are designed for a particular functionality and perform their task more *predictably* than humans, from driving cars to supply chain planning. Frequently, some tasks are not considered AI problems any more, once they are solved without full-fledged intelligence. This phenomenon is known as the "AI effect" [124]. It would be unfair, however, to deny that some current AI systems, especially those that incorporate some learning potential, exhibit some intelligent behaviour.

Anyway, it is not the purpose of this paper to dig further into the time-worn debate between narrow AI vs. general AI. Both approaches are valid and genuine parts of AI research. It is useful to have specialised AI systems that solve specific tasks, as well as systems that have abilities so that they can solve new problems they have never faced before. The intention of stressing this duality is that this should necessarily pervade the evaluation procedures in AI. Specialised AI systems should require a task-oriented evaluation, while general AI systems should require an ability-oriented evaluation.

*This paper corresponds to a lecture given for the Summer School of the Spanish Association for Artificial Intelligence, in A Coruña, Spain, September 2014.

This paper pays attention to the way evaluation is done in AI. As any science and engineering discipline, measuring is crucial for AI. Disciplines progress when they have objective evaluation tools to measure the elements and objects of study, assess the prototypes and artefacts that are being built and examine the discipline as a whole. As we will discuss in subsequent sections, despite the significant progress in the past couple of decades (with the generalisation of several AI benchmarks and competitions) there is still a huge margin of improvement in the way AI systems are evaluated. This is partially because we do not see AI evaluation as a *measurement process*[64]. Also, it is probably a crucial moment to overhaul the way AI evaluation is performed, after the recent progress in areas of AI that are detaching from the narrow AI approach, such as developmental robotics [8], deep learning [7], inductive programming [69, 63][62], artificial general intelligence [60], universal artificial intelligence [88], etc.

By overhauling AI evaluation, we aim at filling a gap, because, to our knowledge, there is no comprehensive analysis about how evaluation is performed in AI and how it can be improved and adapted to the challenges of the future. Some previous works discussing AI evaluation [132, 56, 138, 57, 33, 105, 106, 21, 151, 11, 50, 104, 108, 183, 44, 6, 119, 145] are relatively old, non-comprehensive, restrictive to a specific area of AI, limited to one particular approach and/or focussed on the experimental methodology rather than what is being measured and how. Nonetheless, we will refer to many of these works along the text.

Some ideas of the old analysis still hold today. For instance, in [29] we find criteria for evaluating research problems, methods, implementations, experiments' design, and evaluation of the experiments. In the criteria for experiments' design, we see several of the topics we will address in the paper: "1. How many examples can be demonstrated?" (are they sufficient and qualitative different and illustrative?), "2. Should the program's performance be compared to a standard?", "3. What are the criteria for good performance?", "4. Does the program purport to be general (domain-independent)?" (do the domains being tested constitute a representative class?), and "5. Is a series of related programs being evaluated?". Other statements in [29] are not so up-to-date, and show that there has been an improvement in AI evaluation. For instance, we found the recommendation "that editors, program committees, and reviewers should begin to insist on evaluation.". Today this recommendation has been generalised (e.g., [31] report that more than 60% of ICAIL papers in 1987 did not have any evaluation in front of 20% in 2011). Hence, a lack of evaluation is no longer the problem. However, there is still a great deal of disaggregation, many ad-hoc procedures, bad habits and loopholes about what is being measured and how is being measured. In this paper, the focus will be set on these issues.

We will start with a state of the art of the task-oriented evaluation approach in AI, by far more common in AI research. The notion of performance is relatively easy to determine as it is directly linked to the set or class of problems we are interested in for the evaluation. Nonetheless, we will identify several problems, most of them derived from the confusion of a task definition with its evaluation. An appropriate sampling procedure from the class of problems defining the task is not always easy. We will give some hints to derive better evaluation protocols. With this perspective we will argue that white-box evaluation (by algorithm inspection) is becoming less predominant in AI, and we will focus the rest of the paper to black-box evaluation (by behaviour). We will distinguish three types of behavioural evaluation: by human discrimination (performing a comparison against or by humans), problem benchmarks (a repository or generator of problems) and by peer confrontation (1-vs-1 or multi-agent 'matches'). We will survey some of the competitions and repositories in these three categories and highlight some problems in how these evaluation settings are held and used.

In a second part of the paper, we will pay attention to the more elusive and challenging problem of ability-based evaluation. The three types of evaluation seen for task-oriented evaluation are not directly applicable, as we now do not want to evaluate systems for what they do but for what they are able to (learn to) do. In other words, we are looking for signs or indications that show that the system has a certain ability. One idea that has been around since the inception of AI is to use human (or animal) intelligence tests, such as the IQ-tests used in psychometrics. Each particular test tries to identify a series of exercises that are representative (necessary and sufficient) for a given ability. We will briefly discuss their use and possible adaptation for the evaluation of AI systems. A quite different approach is based on algorithmic information theory (AIT), where problem classes and their difficulty are derived from computational principles. In this way, we are sure about what we are actually evaluating. Also, exercise generators can be derived from first principles.

While task-oriented evaluation is opposed to ability-oriented evaluation in this paper, we can have a more gradual view in terms of task classes that go from specific to general. Also, we will analyse a more unified

view that integrates the different evaluation paradigms and procedures that we find in many disciplines, depending on the subject that is being measured. This view, known as 'universal psychometrics', is based on the notion of 'universal test'. This unified approach makes it possible that the schemas that were identified for the task-based evaluation can be generalised to the ability-based evaluation problem.

The rest of the paper goes along the organisation described above, with two parts: section 2, focussing on task-oriented evaluation and section 3, focussing on ability-oriented evaluation. This is followed by the conclusions, which feature some guidelines about how competitions and problem generators can be improved, integrated or overhauled for a more robust and efficient AI system evaluation.

2 Task-oriented evaluation

AI is a successful discipline. The range of applications has been greatly enlarged over the years. We have successful applications in computer vision, speech recognition, music analysis, machine translation, text summarisation, information retrieval, robotic navigation and interaction, automated vehicles, game playing, prediction, estimation, planning, automated deduction, expert systems, etc. (see, e.g., [139]). Most of these application problems are specific. This implies that the goals are clear and that researchers can focus on the problem. This does not mean that we are not allowed to use more general principles and techniques to solve many of these problems, but that the task is sufficiently specific so that systems can be specialised for these systems. For instance, robotic navigation of a Mars rover can share some of the techniques with a driverless car on Earth, but the final application is extremely specialised in both cases.

This specialisation leads to an application-specific (task-oriented) evaluation. In fact, going from an abstract problem to a specific task is encouraged: "refine the topic to a task", provided it is "representative" [29]. Given a precise definition of the task, we only need to define a notion of performance from it. Clearly, we measure performance, and not intelligence. In fact, many of the most successful AI systems solve each problem in a way that is different to the way humans solve the same problem. Also, AI systems usually include a great amount of built-in programming and knowledge for the task. It is not unfair to say that we evaluate the researchers that have designed the system rather than the system itself. For instance, we can say that it was the research team after Deep Blue [24] (with the help of a powerful computer) who actually defeated Kasparov.

Disregarding who is to praise for each new successful application, AI systems that address specialised problems with a clear performance should be easy to evaluate. The reality is not that straightforward, mostly because there are many different (and usually ad-hoc) evaluation approaches. Let us have a perusal over them.

2.1 Types of performance measurement in AI

An application, as described above, can be characterised by a set of problems, tasks or exercises M. In order to evaluate each exercise $\mu \in M$ we can get a measurement $R(\pi, \mu)$ of the performance of system π. Measurements can be imperfect. Also, the system, the problem or the measurement may be non-deterministic. As a result, it is usual to work with the expected value of the performance of π as $\mathbb{E}[R(\pi, \mu)]$.

The definition of M and R does not specify how we want to aggregate the results when M has more than one problem. The most common approaches are[1]

- Worst-case performance[2]:
$$\Phi_{min}(\pi, M) = \min_{\mu \in M} \mathbb{E}[R(\pi, \mu)]$$
- Best-case performance:
$$\Phi_{max}(\pi, M) = \max_{\mu \in M} \mathbb{E}[R(\pi, \mu)]$$

[1] Worst-case performance and best-case performance are special cases of a rank-based aggregation (using the cumulative distribution of results), with other possibilities such as the median, the first decile, etc. Rank-based aggregation, especially worst-case performance, is more robust to systems getting good scores on many easy problems but doing poorly on the difficult problems.

[2] Note that this formula does not have the size of the instance as a parameter, and hence it is not comparable to the usual view of worst-case analysis of algorithms.

- Average-case performance:

$$\Phi(\pi, M, p) = \sum_{\mu \in M} p(\mu) \cdot \mathbb{E}[R(\pi, \mu)] \quad (1)$$

  where p is a probability distribution on M.

It is assumed that the magnitudes of R for different $\pi \in M$ are commensurate. For instance, if R can range between 0 and 1 for problem μ_1 but ranges between 0 and 10,000 for problem μ_2, the latter will have a much higher weight and will dominate the aggregation. This is not necessarily wrong, e.g., if they are measured with the same unit (e.g., euros). In general, however, R is a construct that needs to be normalised. The choice of a performance metric that is sufficiently normalised such that the results are commensurate is not always easy, but possible to some extent (see, e.g., [183]).

At this point, it is pertinent to make a comment about the well-known no-free-lunch (NFL) theorems [189, 187, 188], as these theorems are usually misunderstood. These theorems state that given all possible problems, *under some particular distributions*, no method can work better than any other on average. The argument to support this interpretation is that, considering all problems, if method π_A is better than method π_B for one problem then π_B will be worse than π_A for another problem. Some people have even interpreted that research in AI (including search and optimisation problems in computer science) is futile. However, the NFL theorems can only be applied when the assumptions hold. The conditions state that M must be infinite and include all possible problems. Also, the problems can be shuffled without affecting the probability, which can be expressed as "block uniformity" [89], for which the uniform distribution would be a special case. Nonetheless, these conditions are not plausible if the problems are taken from the real world. It is unrealistic to assume that the problems we face are taken from a series of random bits, or that a problem, and its opposite problem (whatever it is) are equally probable. Many other distributions are much more plausible. A universal distribution [155, 113], e.g., which is consistent with the idea that problems are generated by physical laws, processes, living creatures, etc., states that random (incompressible) problems are less likely. So, for many distributions p, the conditions of the NFL do not hold and we have that there can be methods π_A and π_B such that: $\Phi(\pi_A, M, p) > \Phi(\pi_B, M, p)$. In fact, there can be optimal methods for inductive inference [107], some free lunches for co-evolution [190], and other areas, although it seems that for optimisation the free lunches are very small [49].

After this clarification, it is relevant to determine how R is going to be obtained. For relatively simple solutions, we can analyse the code or the algorithm of the system π. If the code can be well understood then its computational properties and behaviour can be clearly determined. We use the term 'white-box' evaluation when R is inferred through program inspection or algorithm analysis. White-box evaluation is powerful because we can obtain R theoretically for a given agent π and a problem class M (provided both are defined theoretically). One common type of problems that are evaluated with a white-box approach are those where the solution to the problem has to be correct or optimal (i.e., perfect). In this case, the performance metric R is defined in terms of time and/or space resources. This is the case of classical computational complexity theory. Worst-case analysis is more common than average-case analysis, although the latter has also become popular recently [102, 112, 61]. Nonetheless, many AI problems are so challenging nowadays that perfect solutions are no longer considered as a constraint. Instead, approximate solvers are designed to optimise a performance metric that is defined in terms of the level of *error* of the solution and the time and/or space resources. In this case, the use of an average-case analysis is more common, although worst-case analysis can also be studied under some paradigms (e.g., Probably Approximately Correct learning [166]). In agent theory, the behaviour of the agent (and its properties) can be analysed under some paradigms such as Belief-Desire-Intention (BDI) agents (see, e.g., a testability approach in [186]). The theoretical analysis of 'white-box' evaluation has also been applied to games. For instance, in board games, algorithms can be derived and analysed whether they are optimal, such as noughts and crosses (tic-tac-toe) and English draughts (checkers), the latter solved by Jonathan Schaeffer [141]. Finally, in game theory, the expected pay-off plays the role of R and optimal strategies can be determined for some simple games, as well as equilibria and other properties. In games, some results can be obtained independently of the opponent, but others are only true if we also know the algorithm that the other players are using (so it becomes a double 'white-box' approach to evaluation).

As AI systems become more sophisticated, white-box assessment becomes more difficult, if not impossible, because the unpredictability of complex systems. Many AI systems incorporate many different techniques

and have stochastic behaviours. This is also in agreement with a view of AI as an experimental science [21, 151]. As a result, a black-box approach, is taken[3]. This means that R is obtained exclusively from the behaviour of the system in an empirical way. In this case, average-case evaluation is usual[4].

There are many kinds of black-box assessment in AI, but we can group them into three main categories:

- Human discrimination: The assessment is made by and/or against humans through observation, scrutiny and/or interview. Although it can be based on a questionnaire or a procedure, the assessment is usually informal and subjective. This type of evaluation is common in psychology, ethology and comparative psychology. In AI, this kind of evaluation is not very usual, except for the Turing Test and variants, as we will discuss later on.

- Problem benchmarks: The assessment is performed against a collection or repository of problems (M). This approach is very frequent in AI, where we have problem libraries, repositories, corpora, etc. It is also usual in psychology and comparative psychology, although in these areas the tests are not publicly available to the systems that are to be evaluated. This has been proposed or suggested in AI occasionally (e.g., the "secret generalized methodology" [183]). For instance, M can be generated in real time using a problem generator, which actually defines M and p.

- Peer confrontation: The assessment for (multi-agent) games is performed through a series of (1-vs-1 or n-vs-n) matches. The result is relative to the other participants. Given this relative value, in order to allow for a numerical comparison, sophisticated performance metrics can be derived (e.g., the Elo system in chess [46]).

The combination of some of the above is also common for evaluation. In what follows, we analyse each of the three categories in more detail.

2.2 Evaluation by human discrimination

In this first category we include the evaluation approaches that are performed by a comparison with or by humans. The Turing Test [165, 133] is a case in which there is both comparison against humans and evaluation by human judges. While the 'imitation game' was introduced by Turing as a philosophical instrument in his response to nine objections against machine intelligence, the game has been (mis-)understood as an actual test ever since, with the standard interpretation of one human, one machine pretending to be a human, and a human interrogator through a teletype acting as a judge. The latter must tell which one is the machine and the human.

Not only has the game been taken as an actual test, but it has had several implementations, such as the Loebner Prize[5], held every year since 1991. Despite the criticisms of how this prize is conducted and its interpretation through the years, there have been more implementations. In 2014, Kevin Warwick, professor at the University of Reading organised a similar competition that took place at the Royal Society in London. Even if the results were not significantly different to previous results of the Loebner Prize (or even what Weizenbaum's ELIZA was able to do fifty years ago [180]), the over-reaction and publicity of this outcome were preposterous. The reputation of the implementations of the Turing Test was (further) stained with statements such as this: "If a computer is mistaken for a human more than 30% of the time during a series of five minute keyboard conversations it passes the test. No computer has ever achieved this, until now. Eugene managed to convince 33% of the human judges (30 judges took part [...]) that it was human." [177]. And Warwick goes on: "We are therefore proud to declare that Alan Turing's Test was passed for the first time. [...] This milestone will go down in history as one of the most exciting".

Is the imitation game a valid test? Even assuming that the times and thresholds are stricter than the previous incarnations, the Turing Test has many problems as an intelligence test. First, it is a test of humanity, relative to human characteristics (i.e., anthropocentric). It is neither gradual nor factorial and

[3]The distinction between white and black box can be enriched to consider those problems where the solution must be accompanied by a verification, proof or explanation [66, 3].

[4]Although it is not uncommon, as we will see, that the set of problems from M are chosen by the research team that is evaluating its own method, so the probability to choose from M can be biased in such a way that it is actually a best-case evaluation.

[5]http://www.loebner.net/Prizef/loebner-prize.html.

Evaluation Setting	Description
Loebner Prize[6]	General Turing Test implementation.
U. of Reading TT 2014[7]	General Turing Test implementation.
BotPrize[8]	Contest about bot believability in videogames. [114, 84]
Robo Chat Challenge[9]	Chattering bots competition.
CAPTCHAs[10]	Spotting bots in applications requiring humans. [172, 173]
Humies awards[11]	Human-competitive results using genetic and evolutionary computation. [103]
Graphics Turing Test	Tell between a computer-generated virtual world and a real camera. [126, 16]

Table 1: List of some evaluation settings in the human-discrimination category.

needs human intervention (it cannot be automated). If done properly, it may take too much time. Even so, as we have seen, it can be gamed by non-intelligent chatterbox. As a result, the Turing Test is neither a sufficient nor a necessary condition for intelligence. Despite the criticism, the Turing Test still has many advocates [135]. It is also an inspiration for countless philosophical debates and has led to connections with other concepts in AI or computation [77, 78].

In any case, Turing is not to be blamed by a failure of the Turing Test as a useful test to evaluate AI systems. Turing did not conceive the test as a practical test to measure intelligence up to and beyond human intelligence. He is not to blame for a philosophical construct that has had a great impact in the philosophy and understanding of machine intelligence, but originally a negative impact on its measurement.

Does this mean that we should discard the idea of evaluating AI systems by human judges or by comparing with humans? Not at all. Recently, there have been variants of the Turing Test (Total Turing Tests [146], Visual Turing Tests including sensory information, Toddler Turing Tests [5], robotic interfaces, virtual worlds, etc. [130, 83]) that may be useful for chatterbot evaluation, personal assistants and videogames. It is within the area of videogames where the notion of 'believability' has appeared, which is understood as the property of a bot of looking 'believable' as a human [114, 84]. This term is interesting, as it clearly detaches these tests from the evaluation of intelligence. In videogames, there are applications where we want bots that can fool opponents into thinking that they are just another human player. Other highly subjective properties may also be of interest: enjoyability, resilience, aggressiveness, fun, etc.

Finally, there is a kind of test that is related to the Turing Test, the so-called CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) [172, 173]. It is said to be a 'reverse Turing Test' because the goal is to tell computers and humans apart in order to ensure that an action or access is only performed by a human (e.g., making a post, registering in a service, etc.). CAPTCHAs are quick and practical, omnipresent nowadays. However, they are designed according to the tasks that are solved by the current state of AI technology. At present, for instance, a common CAPTCHA is a series of distorted letters, which are usually easy to recognise by humans but not by machines (e.g., current OCR systems struggle). Logically, when character recognition systems and other techniques improve, current CAPTCHAs are broken (see, e.g., [23]), and CAPTCHAs need to be updated to more distorted words or to other tasks that are beyond AI technology. Similarly, the detection of bots in social networks (sybils) and crowdsourcing platforms rely on tests that are variants of CAPTCHAs, the Turing Test, or the observation and analysis of user profiles and behaviour [27, 27, 176].

Table 1 includes a selection of evaluation settings under the human-discrimination category. As it is not possible to go into the details of all of them because of brevity, let us choose one that is most representative and with a strong future projection: the BotPrize competition, which has been held since 2008. This contest awards the bot that is deemed as more believable (playing like a human) by the other (human) players. The competition uses a first-person shooter videogame, the DeathMatch game type, as used in Unreal Tournament 2004. It is important to clarify that the bots do not process the image but receive a description

[6] http://www.loebner.net/Prizef/loebner-prize.html
[7] http://www.reading.ac.uk/news-and-events/releases/PR583836.aspx
[8] http://botprize.org/
[9] http://www.robochatchallenge.com/
[10] http://www.captcha.net/
[11] http://www.human-competitive.org

of it through textual messages in a specific language through the GameBots2004 interface (Pogamut). For the competition, chatting is disabled (as it is not a chatbot competition). There is a "judging gun" and the human judges also play, trying to play normally (a prize for the judges exists for those that are considered more "human" by other judges).

Some questions have been raised about how well the competition evaluates the believability of the participants. For instance, believability is said to be better assessed from a third-person perspective (judging recorded video of other players without playing) than with a first-person perspective [164]. The reason is that third-person human judges can concentrate on judging and not on not being killed or aiming at high scores. Actually, this third-person perspective is included in the 2014 competition using a crowdsourcing platform [115] so that the 2014 edition incorporates the two judging systems: the First-Person Assessment (FPA), using the BotPrize in-game judging system, and the Third-Person Assessment (TPA), using a crowdsourcing platform. Another issue that could be considered in the future is a richer (and more challenging) representation of the environment, closer to the way humans perceive the images of the game (such as the graphical processing required for the Arcade Learning Environment [12] or the General Video Game Competition [143]).

Finally, as a summary of the limitations and potentials of the human-discrimination category, we first acknowledge that some variants are being useful. However, the format differs significantly from a standard Turing Test. For instance, the human-discrimination approach to evaluation can be just solved by a more traditional interview format with a procedure or storyline (as in psychology interviews), or by an evaluation through observation (using a committee of dedicated judges). This casts doubts about whether evaluation by imitation using the standard interpretation of the Turing Test is practical for task-oriented evaluation in AI. It is the concept that is useful, and deserves being adapted to several applications.

2.3 Evaluation through problem benchmarks

In this very common approach to evaluation, M is defined as a set of problems. This fits equation 1 perfectly. Necessarily, the quality of the evaluations depends on M and how exhaustively this set is explored. There are other issues that could compromise the quality of the measurement. For instance, when M is a *public* problem repository and is not very large, we have that the systems can specialise for M. Also, the solutions may also be available beforehand, or can be inferred by humans, so the systems can embed part of the solutions. In fact, a system can succeed in a benchmark with a small size of M by using a technique known as the "big switch", i.e., the system recognises which problem is facing and uses the hardwired solution for that specific exercise. Things can become worse if the selection of examples from M is made by the researchers themselves (e.g., the usual procedure in machine learning of selecting 10 or 20 datasets from the UCI repository [9], as we will discuss below). In general, the size of M and a bona fide attitude to research somewhat limit these concerns. Nonetheless, it is generally acknowledged that most systems actually embed what the researchers have learnt from M. In a way, these benchmarks actually evaluate the researchers, not their systems.

The above-mentioned problem is known as 'evaluation overfitting' [183], 'method overfitting' problem [50] or "clever methods of overfitting" [104]. To avoid or reduce this problem, it is much better if M is very large or infinite, or at least the problems are not disclosed until evaluation time. Problem generators are an alternative. However, it is not always easy to generate a large M of realistic problems. Generators can be based on the use of some prototypes with parameter variations or distortions. These prototypes can be "based on reality", so that the generator "takes as input a real domain, analyses it automatically and generates deformations [...] that follow certain high-level characteristics" [44]. More powerful and diverse generators can be defined by the use of problem representation languages. A general and elegant approach is to determine a probabilistic or stochastic generator (e.g., a grammar) of problems, which directly defines the probability p for the average-case performance equation 1. Nonetheless, it is not easy to make a generator that can rule out unusable or Frankenstein problems. In other domains, problems are taken from real life (e.g., pedestrian detection), and having a large number of labelled examples is very expensive. Virtual simulators are becoming common to create problems [170].

When the set of problems is large or generated, we clearly cannot evaluate AI systems efficiently with the whole set M. So we need to do some sampling of M. It is at this point when we need to distinguish the benchmark or problem definition from an effective evaluation. Assume we have a limited number of exercises

n that we can administer. The goal will be to reduce the variance of the measurement given n. One naive approach is to sort M by decreasing p and evaluate the system with the first n exercises. This maximises the accumulated mass for p for a given n. One problem about this procedure is that it is highly predictable. Systems will surely specialise on the first n exercises. Also, this approach is not very meaningful when R is non-deterministic and/or not completely reliable. Repeated testing may be necessary, which raises the question of whether to explore a higher n or to perform more repetitions.

Random sampling using p seems to be a more reasonable alternative. As said above, if R is non-deterministic and/or subject to measurement error, then random sampling can be with replacement. If M and p define the benchmark, is probability-proportional sampling on p the best way to evaluate systems? The answer is no, in general. There are better ways of approximating equation 1. The idea is to sample in such a way that the diversity of the selection is increased. This 'diversity-driven sampling" is related to several kinds of sampling, such as importance sampling [156], stratified sampling [28] and other forced Monte Carlo procedures. The key issue is that we use a *different* probability distribution for sampling. Although there are many ways of obtaining a 'diverse' sample, we just highlight two main approaches that can be useful for AI evaluation:

- Information-driven sampling: Assume that we have a similarity function $sim(\mu_1, \mu_2)$, which indicates how similar (or correlated) exercises μ_1 and μ_2 in M are. In this case, we need to sample on M such that the accumulated mass on p is high and that diversity is also high. The rationale is that if μ_1 and μ_2 are very similar, using one of them can 'fill the gap' of the other, and we can assume as if both μ_1 and μ_2 had been explored, actually accumulating $p(\mu_1) + p(\mu_2)$. One possible way of doing this is by cluster sampling. Information-driven sampling suffers from the need of defining the similarity function sim. An alternative is to derive m features that describe the exercises, so creating an m-dimensional space where distances and other topological information can be used to support the notion of diversity (and performing clustering). An example of this procedure is shown in Figure 1 (left).

- Difficulty-driven sampling. A set M can contain very easy and very challenging problems. Using easy problems for good systems or difficult problems for bad systems is not very optimal. The idea to optimise the evaluation is to choose a range of difficulties for which the evaluation results may be informative (or to give higher probability to exercises inside this range), as in Figure 1 (right). This procedure is done to a greater or lesser degree in many evaluations and benchmarks in AI. In fact, more challenging problems are usually added over the years, as the systems are able to solve the easy problems (that soon become 'toy problems'). One of the crucial points of difficulty-driven sampling is the definition of a difficulty function $d : M \to \mathbb{R}^+$. Ideally, we would like that for every $\pi, \Phi(\pi, \mu_1, p) > \Phi(\pi, \mu_2, p)$ iff $d(\mu_1) < d(\mu_2)$. In practice, this condition is too strong, and more flexible characterisations are expected, such as that for every π, and two difficulties a and b such that $a \leq b$ we have that $\Phi(\pi, M_a, p) \geq \Phi(\pi, M_b, p)$ (where M_a denotes all the exercises in M of difficulty a). This could still too strong and we may use a relaxed version such that for every π, there is a t such that for all a and $b \geq a + t$: $\Phi(\pi, M_a, p) \geq \Phi(\pi, M_b, p)$. In experimental sciences, we have a population-based view of difficulty such that $d(\mu)$ is monotonically decreasing on $\mathbb{E}_{\pi \in \Omega}[\Phi(\pi, \mu, p)]$, where Ω is a population of subjects, agents or systems that are evaluated for the same problem class. In fact, Item Response Theory [47] in psychometrics follows this approach. Finally, we can derive the difficulty of a problem as a function of the complexity of the problem itself. The complexity metric can be specific to the application (such as the complexity for mazes in [10, 192] or grid-world domains in [160]) or it can be a more general approach (e.g., Kolmogorov complexity). Note that some of the definitions of difficulty above would not be possible for a set M and distribution p if the conditions of the NFL theorem held.

Both the information-driven sampling and the difficulty-driven sampling can be made adaptive. The first is represented by what is known by adaptive cluster sampling [148] and it is common in population surveys and many experimental sciences. However, when evaluating performance, it is difficulty-driven sampling that has been used more systematically in the past, especially in psychometrics. In psychometrics, difficulty is inferred from a population of subjects (in the case of AI, this could be a set of solvers or algorithms). Instead of difficulty, items are analysed by proficiency, represented by θ, a corresponding concept to difficulty from the point of view of the solver (higher problem difficulty requires higher agent proficiency).

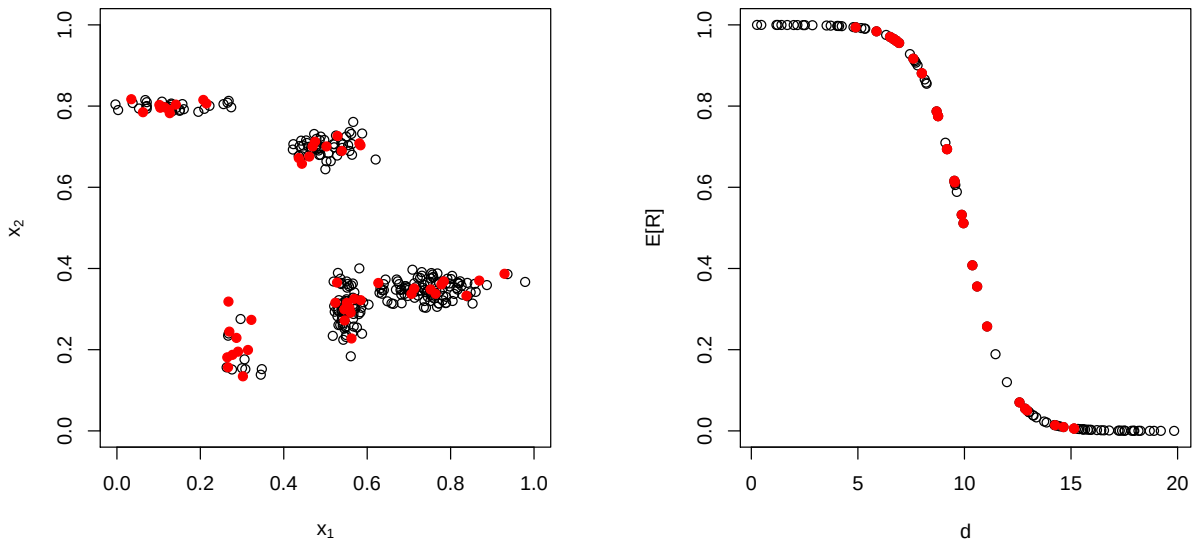

Figure 1: Left: a repository M with $|M| = 300$ exercises shown with empty black circles. Two features x_1 and x_2 are used to describe the most relevant characteristics of the exercises (according to diversity). These features are used to cluster them into five groups. Next, cluster sampling is performed with a sample size of $n = 50$. Clusters are of different size (60, 20, 70, 110, 30) but 10 samples (shown in solid red circles) are taken from each cluster. Because of the constant number of examples per cluster, in order to estimate Φ, measurements for under-represented clusters are multiplied by their size. Right: a repository of $|M| = 100$ exercises. A measure of difficulty d has been derived that is monotonically decreasing with (estimated) expected performance (for a group of agents or for the problem overall). Only $n = 30$ exercises are sampled in the area where the results may be most informative.

Item response theory (IRT) [47] estimates mathematical models to infer the associated probability and informativeness estimations for each item. When R is discrete or bounded, one very common model is the three-parameter logistic model, where the item response function (or curve) corresponds to the probability that an agent with proficiency θ gives a correct response to an item. This model is characterised as follows:

$$p(\theta) \triangleq c + \frac{1-c}{1+e^{-a(\theta-b)}}$$

where a is the *discrimination* (the maximum slope of the curve), b is the *difficulty* or item location (the value of θ leading to a probability half-way between c and 1, i.e., $(1+c)/2$), and c is the chance or asymptotic minimum (the value that is obtained by *random* guess, as in multiple choice items). The zero-ability expected result is given when $\theta = 0$, which is exactly $z = c + \frac{1-c}{1+e^{ab}}$. Figure 2 (left) shows an example of a logistic item response curve.

For continuous R, if they are bounded, the logistic model above may be appropriate. On other occasions, especially if R is unbounded, a linear model may be preferred [127, 51]:

$$X(\theta) \triangleq z + \lambda\theta + \epsilon$$

where z is the intercept (zero-ability expected result), λ is the loading or slope, and ϵ is the measurement error. Again, the slope λ is positively related to most measures of discriminating power [52]. Figure 2 (right) shows an example of a linear item response curve.

Working with item response models is very useful for the design of tests, because if we have a collection of items, we can choose the most suited one for the subject (or population) we want to evaluate. According to the results that the subject has obtained on previous items, we may choose more difficult items if the subject has succeeded on the easy ones, we may look for those items that are most discriminating (i.e., most informative) in the area we have doubts, etc. Note that discrimination is not a global issue: a curve may have a very high slope at a given point, so it is highly discriminating in this area, but the curve will almost

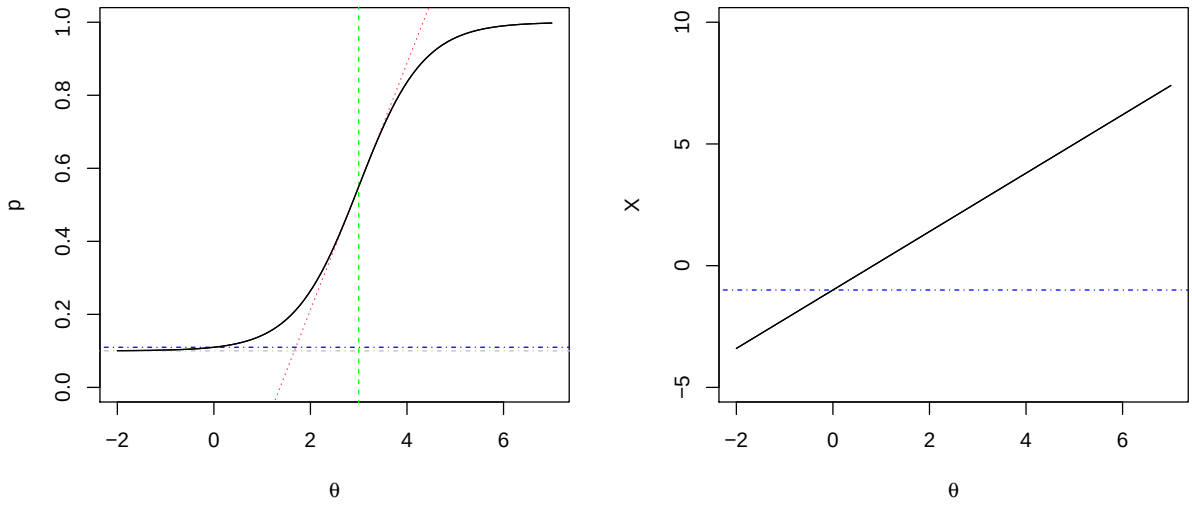

Figure 2: Left: item response function (or curve) for a binary score item with the following parameters for the logistic model: discrimination $a = 1.5$, item location $b = 3$, and chance $c = 0.1$. The discrimination is shown by the slope of the curve at the midpoint: $a(1 - c)/4$ (in dotted red), the location is given by b (in dashed green) and the chance is given by the horizontal line at c (in dashed-dotted grey), which is very close to the zero-proficiency expected result $p(\theta) = z$ (here at 0.11). Right: A linear model for a continuous score item with parameter $z = -1$ and $\lambda = 1.2$. The dashed-dotted line shows the zero-ability expected result.

be flat when we are far from this point. Conversely, if we have a low slope, then the item covers a wide range of difficulties but the result of the item will not be so informative as for a higher slope.

Figure 3 shows an example of an adaptive test using IRT. The sequence of exercise difficulties is shown on the left. The plot on the right shows that averaging the results (especially here, as the outcome of R is discrete, either 0 or 1) makes the estimation of Φ more difficult with a non-adaptive test.

Table 2 includes a selection of evaluation settings in the problem benchmarks category. We see the variety of repositories, challenges and competitions. As it is impossible to survey all of them in detail, we

[12] http://www.cs.miami.edu/~tptp/CASC/
[13] http://termination-portal.org/wiki/Termination_Competition_2014
[14] http://www.rl-competition.org/
[15] http://www.sygus.org/
[16] http://www.aerialroboticscompetition.org/
[17] http://archive.darpa.mil/grandchallenge04/index.htm
[18] http://archive.darpa.mil/grandchallenge/
[19] http://www.darpa.mil/cybergrandchallenge/
[20] http://www.theroboticschallenge.org/
[21] http://ipc.icaps-conference.org/
[22] http://archive.ics.uci.edu/ml/
[23] http://sci2s.ugr.es/keel/datasets.php
[24] http://prtools.org/
[25] http://www.sigkdd.org/kddcup/index.php
[26] http://www.kaggle.com/
[27] http://pan.webis.de/
[28] http://www.gvgai.net/
[29] http://prize.hutter1.net/
[30] http://mattmahoney.net/dc/text.html
[31] http://www.statmt.org/europarl/
[32] http://www.statmt.org/setimes/
[33] http://matrix.statmt.org/matrix/info
[34] https://www.ldc.upenn.edu/new-corpora
[35] http://www.arcadelearningenvironment.org/
[36] http://gpbenchmarks.org/
[37] http://www.movingai.com/benchmarks/
[38] http://www.fira.net/contents/sub03/sub03_1.asp

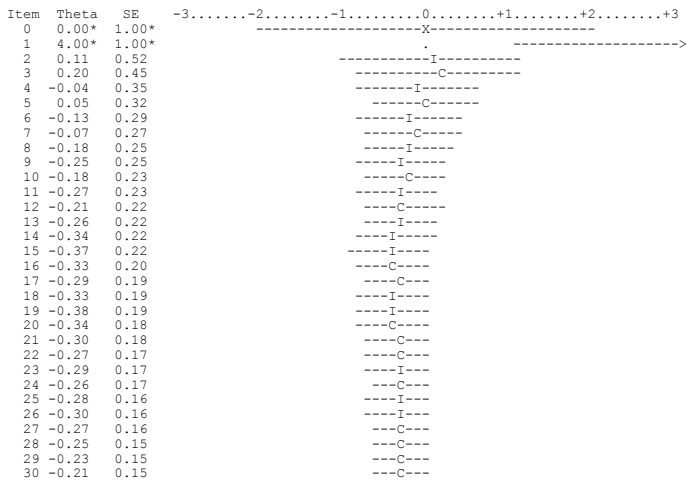

```
Item  Theta   SE    -3.......-2........-1.........0........+1........+2........+3
  0   0.00*  1.00*             --------------------X--------------------
  1   4.00*  1.00*                                 .          --------------------->
  2   0.11   0.52                          ----------I----------
  3   0.20   0.45                            ----------C----------
  4  -0.04   0.35                          -------I-------
  5   0.05   0.32                            ------C------
  6  -0.13   0.29                          ------I------
  7  -0.07   0.27                           ------C-----
  8  -0.18   0.25                          -----I-----
  9  -0.25   0.25                         -----I-----
 10  -0.18   0.23                          -----C----
 11  -0.27   0.23                         -----I----
 12  -0.21   0.22                          ----C-----
 13  -0.26   0.22                          ----I----
 14  -0.34   0.22                         ----I-----
 15  -0.37   0.22                        -----I----
 16  -0.33   0.20                         ----C----
 17  -0.29   0.19                          ----C---
 18  -0.33   0.19                         ----I----
 19  -0.38   0.19                         ----I----
 20  -0.34   0.18                         ----C----
 21  -0.30   0.18                          ----C---
 22  -0.27   0.17                          ----C---
 23  -0.29   0.17                          ----I---
 24  -0.26   0.17                           ---C---
 25  -0.28   0.16                         ----I---
 26  -0.30   0.16                         ----I---
 27  -0.27   0.16                          ---C---
 28  -0.25   0.15                          ---C---
 29  -0.23   0.15                          ---C---
 30  -0.21   0.15                          ---C---
```

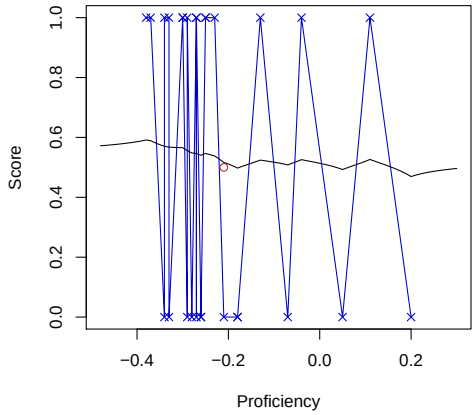

Figure 3: An example of an IRT-based adaptive test (freely adapted from [179, Fig. 8]). Left: the process and proficiencies (*thetas*) used until convergence. The final proficiency calculated by the test was −0.21 with a standard error of 0.15. Right: The results shown on a plot. The black curve shows a Euclidean kernel smoothing with a constant of 0.1.

Evaluation Setting	Description
CADE ATP System Competition[12]	Theorem proving [162] using the TPTP library [161].
Termination Competition[13]	Termination of term rewriting and programs [120].
The reinforcement learning competition[14]	Reinforcement learning [184].
Syntax-guided synthesis competition[15]	Program synthesis [4].
International Aerial Robotics Competition[16]	Pilotless aircraft competition.
DARPA Grand Challenge[17]	Autonomous ground vehicles.
DARPA Urban Challenge[18]	Driverless vehicles.
DARPA Cyber Grand Challenge[19]	Computer security.
DARPA Save the day[20]	Rescue Robotic challenge [95].
The planning competition[21]	Planning [116].
UCI[22] and KEEL[23]	Machine learning dataset repositories [9] [1].
PRTools[24]	Pattern recognition problem repository.
KDD-cup challenges[25] and kaggle[26]	Machine learning and data mining competitions.
Plagiarism detection[27]	Plagiarism detection, authorship and social software misuse [134].
The General Video Game Competition[28]	General video game players [143].
Hutter Prize[29] and related benchmarks[30]	Text compression.
Pedestrian benchmarks	Pedestrian detection [59]
Europarl[31], SE times corpus[32], the euromatrix[33]	Machine translation corpora [157].
Linguistic data consortium corpora[34]	NLP corpora.
The Arcade Learning Environment[35]	Atari 2600 videogames (reinforcement learning) [12].
GP benchmarks[36]	Genetic programming [125, 182].
Pathfinding benchmarks[37]	Gridworld domains (mazes) [160].
FIRA HuroCup[38]	Humanoid robot competitions [6]

Table 2: List of some evaluation settings in the problem-benchmarks category.

will focus on one of them, perhaps the most widespread repository in computer science, the UCI machine learning repository [9]. Most of the discussion below is applicable to other repositories and, to some extent, to competitions and challenges in machine learning.

The UCI repository includes many supervised (classification and regression) and some unsupervised datasets. The repository is publicly available and is regularly used in machine learning research. The usage procedure, which is referred as "The UCI test" [118] or the "de facto approach" [44][96], follows the general form of equation 1 where M is the repository, p is the choice of datasets and R is one particular performance metric (accuracy, AUC, Brier score, F-measure, MSE, etc. [53, 76]). With the chosen datasets, several algorithms (where one or more are usually introduced by the authors of the research work) can be evaluated by their performance on the datasets. The aggregation over several datasets according equation 1, however, is not very common in machine learning, as there is the general belief that averaging the results for several datasets is wrong, as results are not commensurate (see, e.g., [34]). We already discussed this issue in section 2.1 and saw that there are ways to normalise the performance metric or use some utility measure instead (e.g., what are the costs, in euros, of false positives and false negatives for each dataset) such that they can be aggregated. Nonetheless, statistical tests are the predominant and encouraged approach to evaluation validation by the machine learning research community.

"The UCI test" can be seen as a bona-fide mix of the problem benchmark approach and the peer confrontation approach. Even if there is a repository (M), only a few problems are chosen, and can be cherry-picked (p is changing and arbitrary). Also, as the researchers' algorithm has to be compared with other algorithms, a few competitors are chosen, which can also be cherry-picked, without much effort on fine-tuning their best parameters. Finally, as the results are analysed by statistical tests, cross-validation or other repetition approaches are used to reduce the variance of $R(\pi, \mu, p)$ so that we have fewer "ties". This procedure frequently leads to claims about new methods being better than the rest. Many of these claims are, apart from uninteresting, dubious, even for papers in good venues. Nonetheless, the UCI repository is not to blame for this procedure, but a methodology where statistical significance for a few datasets is more valued than a commensurate average aggregate performance on a large collection of datasets.

As a result, there have been suggestions of a better use of the UCI repository. These suggestions imply an improvement of the procedure but also of the repository itself. For instance, UCI+, "a mindful UCI" [118], proposes the characterisation of the problems in the UCI repository by a set of complexity measures from [85]. This characterisation can be used to make samples that are more diverse and representative. Also, they discuss the notion of a problem being 'challenging', trying to infer a notion of 'difficulty'. In the end, an artificial dataset generator is proposed to complement the original UCI dataset. It is a distortion-based generator (similar to Soares's UCI++ [153]). Finally, [118] suggest ideas about sharing and arranging the results of previous evaluations so that each new algorithm can be compared immediately with many other algorithms using the same experimental setting. This idea of 'experiment database' [168] has already been set up. Openml[39] [167, 169], is an open science platform that integrates machine learning data, software and results An automated submission procedure, such as Kaggle, if performed for a wide range of problems at a time, could be a way of controlling some of the methodological problems of how the UCI repository is used.

Although some of these improvements are in the line of better sampling approaches (more representative and more effective), there are still many issues about the way these repositories are constructed and used. The complexity measures could be used to derive how representative a problem is with respect to the whole distribution in order to make a more adequate sampling procedure (e.g., a clustering sampling). Also, a pattern-based generator instead of a distortion-based generator could give more control of what is generated and its difficulty. This could be done with a stochastic generative grammar for different kinds of patterns, as is usually done with artificial datasets, using Gaussians or geometrical constructs. Finally, if results are aggregated according to equation 1, the experimental setting and the use of repetitions should be overhauled. For instance, by using 20 different problems with 10 repetitions using cross-validation (a very common setting in machine learning experiments) we have less information than by using 200 different problems with 1 repetition. Choosing the least informative procedure only makes sense because of the way results are fitted into the statistical tests and also because repetitions usually involve less effort than preparing a large number of datasets.

Overall, even if the UCI repository and machine learning are very particular, many of the benchmarks

[39] http://openml.org/.

in Table 2 suffer from the same problems about how representative the problems are (if M is small) or how representative the sample is (if M is large). Other problems are the estimation of task difficulty and whether M is able to discriminate between a set of AI systems. Also, none of the benchmarks in AI is adaptive.

2.4 Evaluation by peer confrontation

In the evaluation by peer confrontation, we evaluate a system by confronting it to another system. This usually means that a match is played between peers. This is usual for games (including game theory) and part of multi-agent research. The results of each match (possibly repeated with the same peer) may serve as an estimation of which of the two systems is best (and how much). Nonetheless, the main problem about this approach is that the results are relative to the opponents. This is natural in games, as people are said to be good or bad at chess, for instance, depending of whom they are compared with.

Despite this relative character of the evaluation, we can still see the average performance according to equation 1. In order to do this, we must first identify the set of opponents Ω. Then, the set of problems M is enriched (or even substituted) by the parametrisation of each single game (e.g., chess) with different competitors from Ω. In 1-vs-1 matches we have that $|M| = |\Omega| - 1$ (if we do not consider a match between a system and itself). In other multi-agent situations where many agents play at the same time, $|M|$ can grow combinatorially on $|\Omega|$.

Nonetheless, for AI research, our main concern is about robustness and standardisation of results. For instance, how can we compare results between two different competitions if opponents are different? If these competitions are performed year after year, how can we compare progress? If there are common players, we can use rankings, such as the Elo ranking [46], or more sophisticated rating systems [152, 121], to see whether there are progress. In fact, it would be very informative for AI competitions based on peer confrontation to keep all participants from previous editions in subsequent editions. However, this comes with a drawback, as systems could specialise to the kind of opponents that are expected in a competition. If a high percentage of competitors are inherited from previous editions, specialisation to those old (and bad) systems could be common. It is insightful to think how many of these issues are addressed in sport competitions. For instance, some tournaments adapt their matches according to previous information (by round, by ranking, etc.). In fact, a league may be redundant (for the same reasons why the information-driven or difficulty-driven sampling are introduced) and other tournament arrangements are more effective with almost the same robustness and much fewer matches.

As an alternative, games and multi-agent environments could be evaluated against standardised opponents. However, how can we choose a set of standardised opponents? If the opponents are known, the systems can be specialised to the opponents. For instance, in an English draughts (checkers) competition, we could have players being specialised to play against Chinook, the proven optimal player [141]. Again, this ends up again in the design of an opponent generator. This of course does not mean a random player (which is usually very bad), but players that can play well. One option is to use old systems where some parameters are changed. Alternatively, a more far-reading approach is to define an agent language and generate players (programs) with that language. As it is expected that this generation will not achieve very good players (otherwise we would be facing a very simple problem), a possible solution is to give more information and resources to these standardised opponents to make them more competitive (e.g., in some applications these opponents could have more sophisticated sensor mechanisms or some extra information about the match that regular players do not have).

Be the set Ω composed of old opponents or generated opponents, we need to assess whether Ω is sufficiently challenging and whether it is able to discriminate the participants. For instance, some competitions in AI finally award a champion, but there is the feeling that the result is mostly arbitrary and caused by luck, as happens with many sport competitions[40]. How can we assess whether the set Ω has sufficiently difficulty and discriminating power? This is of course a hard problem, which has recently be analysed in [70]. For instance, Figure 4 shows the distribution of results of an agent competing in a multi-agent system according to the complexity of the agent. The difficulty and discriminating power varies depending on the opponents (left vs right plots).

Table 3 shows a sample of evaluation settings based on peer confrontation. Once again, because of obvious space constraints, we will just choose one representative and interesting case from the table. We will

[40] Statistical tests are not used to determine when a contestant can be said to be significantly better than another.

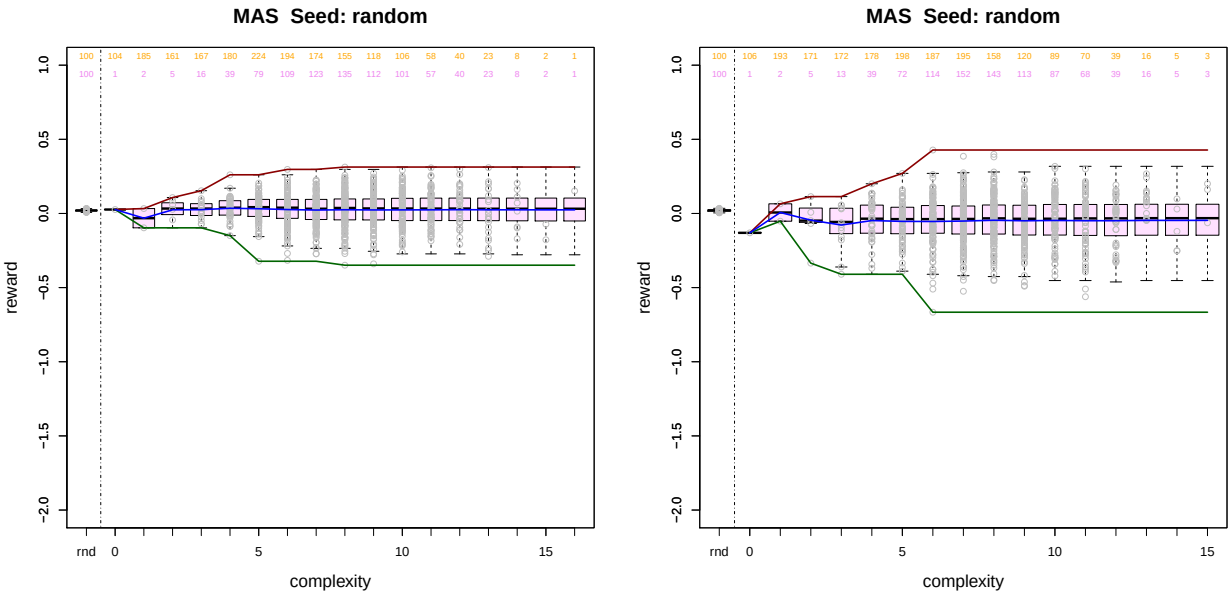

Figure 4: We show the distributions of reward (roughly corresponding to R in this paper) for different configurations for the multi-agent system SCMAS introduced in [70]. Left: the plot shows the results when we confront each of the 2,000 policies with 50 different teams of competitors (with different seeds for the generator also). This means that we have $2{,}000 \times 50 = 100{,}000$ experiments (300 environment steps each). The results for a random agent (rnd) are also shown for comparison. Right: results when we choose the best 8 agents from the previous experiment. We see a wider range of results (but note that the average reward is lower).

Evaluation Setting	Description
Robocup[42] and FIRA[43]	Robotics (robot football/soccer) [100, 99].
General game playing AAAI competition[44]	General game playing using GDL [58].
World Computer Chess Championship[45]	Chess.
Computer Olympiad[46]	Board games.
Annual Computer Poker Competition[47]	Poker.
Trading Agents Competition[48]	Trading agents [181, 98].
Warlight AI Challenge[49]	Strategy games (Warlight).

Table 3: List of some evaluation settings in the peer-confrontation category.

discuss the General Game Competition, which has been run yearly since 2005. According to the webpage[41], "general game players are systems able to accept descriptions of arbitrary games at runtime and able to use such descriptions to play those games effectively without human intervention. In other words, they do not know the rules until the games start". Games are described in the language GDL (Game Description Language). The description of the game is given to the players. Different kinds of games are allowed, such as noughts and crosses (tic tac toe), chess, in static or dynamic worlds, with complete or partial information, with varying number of players, with simultaneous or alternating plays, etc. The competition consists of several rounds, qualifications, etc. For the competition, games are chosen —non-randomly, i.e., manually by the organisers— from the pool of games already described in GDL and new games are also newly introduced for the competition. As a result, game specialisation is difficult.

Despite being one of the most interesting AI competitions, there is still some margin for improvement. For instance, a more sophisticated analysis of how difficult and representative each problem is would be useful. For instance, several properties about the adequacy of an environment or game for peer-confrontation evaluation could be identified and analysed depending on the population of opponents that is being considered [93]. Also, rankings (e.g., using the Elo system mentioned above) could be calculated, and former participants could be kept for the following competitions, so there are more participants (and more overlap between competitions). A more radical change would be to learn without the description of the game, as a reinforcement learning problem (where the system learns the rules from many matches). An adaptation between the general game playing and RL-glue, which is used in the reinforcement learning competition, to make this possible has been done in [13].

Summing up our observations on peer confrontation problems, we see that the dependency on the set $|\Omega|$ makes this kind of evaluation more problematic. Nonetheless, as AI research is becoming more socially oriented, with significantly more presence of multi-agent systems and game theory, an effort has to be done to make this kind of evaluation more systematic, instead of the plethora of arrangements that we see in sports, for instance.

As a summary of this whole section about task-oriented evaluation, we have identified many issues in many evaluation settings in AI. Nonetheless, the three types of evaluation settings have their niches of application and task-specific evaluation is the right one in applications such as engineering, medicine, military devices, etc. In fact, the series of workshops on Performance Metrics for Intelligent Systems, held since 2000 at the National Institute of Standards & Technology [48, 128, 119, 145] is a good example of the usefulness of this kind of evaluation. However, we hope that some ideas above can be used to make the evaluation more controlled, automated and robust.

3 Towards ability-oriented evaluation

Many areas AI is successful nowadays took a long time to flourish in applications (e.g., driverless cars, machine translators, game bots, etc.). Most of them correspond to specific tasks and require task-oriented evaluation. Other tasks that are still not solved by AI technology are already evaluated in this way and will be successful one day. However, if instead of AI applications we think about AI systems, we see that there are some kinds of AI systems for which task-oriented evaluation is not appropriate. For instance, cognitive robots, artificial pets, assistants, avatars, smartbots, smart houses, etc., are not designed to cover one particular application but are expected to be customised by the user for a variety of tasks. In order to cover this wide range of (previously unseen) tasks, these systems must have some abilities such as reasoning skills, inductive learning abilities, verbal abilities, motion abilities, etc. Hence, this entails that apart from task-oriented evaluation methods we may also need ability-oriented evaluation techniques.

Things are more conspicuous when we look at the evaluation of the *progress* of AI as a discipline. If we look at AI with Minsky's 1968 definition seen in the introduction, i.e., by achievement of tasks that would require intelligence, AI has progressed very significantly. For instance, one way of evaluating AI progress is to look at a task and check in which category an AI system is placed: *optimal* if no other system can perform better, *strong super-human* if it performs better than all humans, *super-human* if it performs better than most humans, *par-human* if it performs similarly to most humans, and *sub-human* if it performs worse than most humans [137]. Note that this approach does not imply that the task is necessarily evaluated with a human-discriminative approach. Having these categories in mind, we can see how AI has scaled up for many tasks, even before AI had a name. For instance, calculation became super-human in the nineteenth century, cryptography in the 1940s, simple games such as noughts and crosses became optimal in 1960s, more complex games (draughts, bridge) a couple of decades later, printed (non-distorted) character recognition

[41] http://games.stanford.edu/
[42] http://www.robocup.org/
[43] http://www.fira.net
[44] http://games.stanford.edu/
[45] http://www.icga.org/
[46] http://www.icga.org/
[47] http://www.computerpokercompetition.org/
[48] http://tradingagents.eecs.umich.edu/
[49] http://theaigames.com/competitions/warlight-ai-challenge/rules

in the 1970s, statistical inference in the 1990s, chess in the 1990s, speech recognition in the 2000s, and TV quizzes, driving a car, technical translation, Texas hold 'em poker in the 2010s. According to this evolution, the progress of AI has been impressive [18]. The use of human intelligence as a baseline has been used in competitions (such as the humies awards[50]) or to define ratios, where median human performance is set at a zero scale, such as the so-called Turing-ratio [122, 121], with values greater than 0 for super-human performance and values lower than 0 for sub-human performance.

However, let us first realise that no system can do (or can learn to do) all of these things together. The big-switch approach may be useful for a few of them (e.g., a robot with an advanced computer vision system that detects whether it is facing a chess board or a bridge table and then switch to the appropriate program to play the game that it has just recognised). Second, if we look at AI with McCarthy's definition seen in the introduction, i.e., by making intelligent machines, things are less enthusiastic. Not only has the progress been more limited, but also there is a huge controversy for quantifying this progress (in fact, some argue that machines are more intelligent today than fifty years ago while others say that there has been no progress at all other than computational power). Hence, worse than having a poor progress or no progress at all, we regard with contempt that we do not have effective evaluation mechanisms to evaluate this progress. It seems that none of the evaluation settings seen in the previous section is able to evaluate whether the AI systems of today are more intelligent than the AI systems of yore. Also, for developmental robotics and other areas of AI where systems are supposed to improve their performance with time, we want to know if a 6-month-old robot has progressed over its initial state, in the same way that we see how abilities increase and crystallise with humans, from toddlers to adults. Ability-oriented evaluation, and not task-oriented evaluation, seems to have a better chance of answering this question.

To make the point unequivocal, we could even go beyond McCarthy's definition of AI without the use of 'intelligence' and define this view of AI as the *science and engineering of making machines do tasks they have never seen and have not been prepared for beforehand*. Clearly, this view puts more emphasis on learning, but it also makes it crystal clear that task-oriented evaluation, as have been performed for years, would not fit the above definition.

It would be unfair to forget to acknowledge that some attempts seen in the previous section have made an effort for a more general AI evaluation. The general game competition seen in the previous section is one example of how some things are changing in evaluation. Users and researchers are becoming tired of a big-switch approach. They yearn for and conceive systems that are able to cover more and more general task classes. Nonetheless, it is still a limited generalisation, which is too based on a very specific range of tasks. Many good players at the General Game Competition would be helpless at any game of the Arcade Learning Environments, and vice versa. Actually, only some reinforcement learning (and perhaps genetic programming) systems can at least participate in (adaptations to) of both games —excelling in them would not be possible though without an important degree of specialisation.

In the rest of this section we will introduce what an ability is and how they can be evaluated in AI. The title of this section (starting with 'Towards') suggests that what follows is more interdisciplinary and contains proposals that are not well consolidated yet, or that may even go in the wrong direction. Nonetheless, let us be more lenient and have in mind that ability-based evaluation is much more challenging than task-specific evaluation.

3.1 What is an ability?

We must first clarify that we are talking about *cognitive* abilities, in the same way that in the previous section we referred to *cognitive* tasks. Some AI applications require physical abilities, most especially in robotics, but AI deals with how the sensors and actuators are controlled, not about their strength, consumption, etc. After this clarification, we can define a cognitive ability as *a property of individuals that allows them to perform well in a range of information-processing tasks*. At first sight this definition may just look like a change of perspective (from problems to systems). However, what we see now is that the ability is required, and performance is worse without featuring the ability. In other words, the ability is necessary but it does not have to be sufficient (e.g., spatial abilities are necessary but not sufficient for driving a car). Also, the ability is assumed to be general, to cover a range of tasks. Actually, general intelligence would be one of

[50] www.human-competitive.org.

these cognitive abilities, one that covers *all* cognitive tasks: "general intelligence is a very broad trait that encompasses quickness and quality of response to all cognitive tasks" [159].

The major issue about abilities is that they are 'properties', and as such they have to be conceptualised and identified. While tasks can be seen as measuring *instruments*, abilities are *constructs*. In psychology, many different cognitive abilities have been identified and have been arranged in different ways [142]. For instance, one well-known comprehensive theory of human cognitive abilities is the Cattell-Horn-Carroll theory [97]. Figure 5 shows a graphical representation of these abilities. The top level represents the g factor or general intelligence, the middle level identifies a set of broad abilities and the bottom level may include many narrow abilities. Again, this top level seems to saturate all tasks: "g is common to all cognitive tasks including learning tasks" [2].

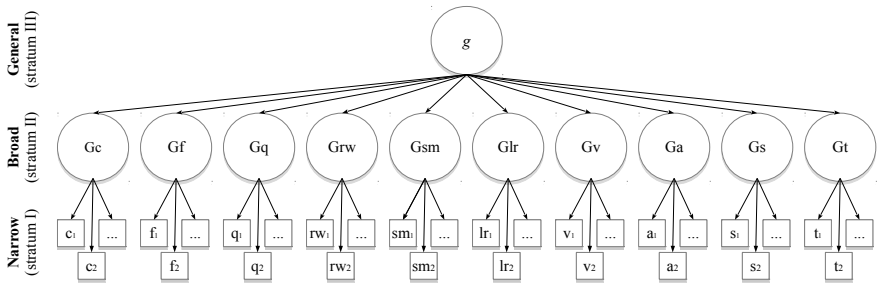

Figure 5: Cattell-Horn-Carroll's three stratum model. The broad abilities are Crystallised Intelligence (Gc), Fluid Intelligence (Gf), Quantitative Reasoning (Gq), Reading and Writing Ability (Grw), Short-Term Memory (Gsm), Long-Term Storage and Retrieval (Glr), Visual Processing (Gv), Auditory Processing (Ga), Processing Speed (Gs) and Decision/Reaction Time/Speed (Gt).

Interestingly, this is not surprising from an AI standpoint. The broad abilities seem to correspond to subfields in AI. For instance, looking at any AI textbook (e.g., [139]), we can enumerate areas such as problem solving, use of knowledge, reasoning, learning, perception, natural language processing, etc., that would roughly correspond to some of the cognitive abilities in Figure 5.

Can we evaluate broad abilities as we did for specific tasks? Application-specific (task-oriented) approaches will not do. But is ability-oriented evaluation ready for this? The answer, as we will see below, is that this type of evaluation is still in a very incipient stage in AI. There are several reasons for this. First, general (ability-oriented) evaluation is more challenging. Second, we no longer have a clear definition of the task(s). In fact, defining the ability depends on a conceptualisation, and from there we need to find a set of representative exercises that require the ability. And third, there have not been too many general AI systems to date, so task-oriented evaluation has seemed sufficient for the evaluation of AI systems so far. However, things are changing as new kinds of AI systems (e.g., developmental robotics) are becoming more general.

Before starting with some approaches in the direction of ability-oriented evaluation, it can be argued that some existing evaluation settings in AI are already ability-oriented. For example, even if the planning competition features a set of tasks, it goes around the *ability* of planning, which is more general than any particular task. However, the systems are not able to determine when planning is required for a range of problems. In other words, the ability is not a resource of the system, but the very goal of the system. In the end, it is the researchers who incorporate planning modules in several application-specific systems, and not the systems that independently enable their planning abilities to solve a new problem.

3.2 The anthropocentric approach: psychometrics

Psychometrics was developed by Galton, Binet, Spearman and many others at the end of the XIXth century and first half of the XXth century. An early concept that arose was the need of distinguishing tasks requiring very specific knowledge or skills from general abilities. For instance, an "idiot savant" could have a lot of knowledge or could have developed a sophisticated skill during the years for some specific domain, but could be obtuse for other problems. On the contrary a very able person with no previous knowledge could perform well in a range of tasks, provided they are culture-fair. This distinction took several decades to consolidate.

In a way, this bears resemblance with the narrow vs. general dilemma in AI.

Psychometrics is concerned about measuring cognitive abilities, personality traits and other psychological properties [158]. Factors differ from abilities, in principle, in that they are obtained through testing and further analysed through systematic approaches based on factor analysis. Some factors have been equated and named after existing abilities while others are 'discovered' and receive new technical names. Several indices can be derived from a battery of tests by aggregating abilities and factors. One joint index that is usually determined from some of these tests is known as IQ (Intelligence Quotient). Although IQ was originally normalised by the subject's age (hence its name), its value for adults today is normalised relative to an adult population, assuming a normal distribution with mean μ=100 and standard deviation σ=15. This corresponds to a more sophisticated (normalised) aggregation of results for several items, which again resembles our equation 1.

IQ tests incorporate items of variable difficulty. Item difficulty is determined by the percentage of subjects that are able to solve the item, using functional models in Item Response Theory [117, 47], as seen in the previous section. Note that this difficulty assessment is relative to the population and not derived from the nature of the item itself.

IQ tests are easy to administer, fast and accurate, and they are used by companies and governments, essential in education and pedagogy. IQ tests are generally culture-fair through the use of abstract exercises (except for the verbal comprehension abilities).

As they work reasonably well for humans, their use for evaluating machines has been suggested many times, even since the early days of AI, with the goal of constructing "a single program that would take a standard intelligence test" [131]. More recently, their use has been vindicated by Bringsjord and Schmimanski [20, 19], under the so-called 'Psychometric AI' (PAI), as "the *field* devoted to building information-processing entities capable of at least solid performance on all established, validated tests of intelligence and mental ability, a class of tests that includes not just the rather restrictive IQ tests, but also tests of artistic and literary creativity, mechanical ability, and so on". It is important to clarify that PAI is a redefinition or new roadmap for AI —not an evaluation methodology— and does not further develop or adapt IQ tests for AI systems. In fact, PAI does not explicitly claim that IQ tests (or other psychometric tests) are the best way to evaluate AI systems, but it is said that an "agent is intelligent if and only if it excels at all established, validated tests of intelligence" (later broadened to any other psychometric test) [20, 19]. The question of whether these tests are a necessary and sufficient condition for machines and the limitations of PAI as a guide for AI research have been recently discussed in [14].

Not surprisingly, this claim that IQ tests are the best way to evaluate AI system has recently come from human intelligence research. Detterman, editor of the Intelligence Journal, wrote an editorial [35] where he suggested that Watson (the then recent winner of the *Jeopardy!* TV quiz [54]) should be evaluated with IQ tests. The challenge is very explicit: "I, the editorial board of *Intelligence*, and members of the International Society for Intelligence Research will develop a unique battery of intelligence tests that would be administered to that computer and would result in an actual IQ score" [35]. Detterman established two levels for the challenge, a first level, where the type of IQ tests can be seen beforehand by the AI system programmer, and a second level, where the types of tests would have not seen beforehand. Only computers passing the second level "could be said to be truly intelligent" [35]. The need for two levels seems related to the big-switch approach and the problem overfitting issue, which we have already mentioned in previous sections for AI evaluation settings. It is apposite at this point to recall that academic and professional IQ tests and many other standardised psychological tests are never made public, because otherwise people could practise on them and game the evaluation. Note that the non-disclosure of the tests until evaluation time is something that we only find in very few evaluation settings in the previous section.

Detterman was unaware that almost a decade before, in 2003, Sanghi and Dowe [140] implemented a small program (less than 1,000 lines of code) which could score relatively well on many IQ tests, as shown in Table 4. The program used a big-switch approach and was programmed to some specific kinds of IQ tests the authors had seen beforehand. The authors still made the point unequivocally: this program is not intelligent and can pass IQ tests.

While it must be conceded that the results only reach the first level of Detterman's challenge —so there is a test administration issue (i.e., an evaluation flaw)—, there are some weaknesses about human IQ tests that would also arise if a system passed the second level as well. In particular, "the editorial board of *Intelligence*, and members of the International Society for Intelligence Research" could be tempted to devise

Test	IQ Score	Human Average
A.C.E. IQ Test	108	100
Eysenck Test 1	107.5	90-110
Eysenck Test 2	107.5	90-110
Eysenck Test 3	101	90-110
Eysenck Test 4	103.25	90-110
Eysenck Test 5	107.5	90-110
Eysenck Test 6	95	90-110
Eysenck Test 7	112.5	90-110
Eysenck Test 8	110	90-110
IQ Test Labs	59	80-120
Testedich IQ Test	84	100
IQ Test from Norway	60	100
Average	96.27	92-108

Table 4: Results by a rudimentary program for passing IQ tests (from [140]).

or choose those IQ tests that are more 'machine-unfriendly'. If AI systems eventually passed some of them, the battery could be refined again and again, in a similar way as how CAPTCHAs are updated when they become obsolete. In other words, this selection (or battery) of IQ tests would need to be changed and made more elaborate year after year as AI technology advances. Also, the limitations of this approach if AI systems ever become more intelligent than humans are notorious.

The main problem about IQ tests is that they are anthropocentric, i.e., they have been devised for humans and take many things for granted. For instance, most assume that the subject can understand natural language to read the instructions of the exercise. On top of that, they are specialised to some human groups. For instance, tests are significantly different when evaluating small children, people with disabilities, etc. Also, the relation between items and abilities have been studied during the past century exclusively using humans, so it is not clear that a set of items would measure the same ability for a human or for a machine. For instance, is it reasonable to expect that well-established tests of choice reaction time be correlated with intelligence in machines as they are correlated in humans [32]? Or, what makes a set of psychometric tests different from a set of "human intelligence tasks" in Amazon Mechanical Turk [22]. For a more complete discussion about why IQ tests are not ready for AI evaluation, the reader is referred to a response [41] to Detterman's editorial. Having said all this and despite the limitations of IQ tests for AI evaluation, their use is becoming more popular in the past decade (including robotics, [144]) and systems whose results are like those of Table 4 are becoming common (for a survey, see [79], for an open library of IQ tests, see `pebl.sourceforge.net/battery.html`).

As just said, one of the problems of IQ tests is that they are specialised for humans. In fact, standardised adult IQ tests do not work with people with disabilities or children of different ages. In a similar way, we do not expect animals to behave well on a standard human IQ test, starting from the fact that they will not be able to read the text. This leads us to the consideration of how cognitive abilities are evaluated in animals. Comparative psychology and comparative cognition [149, 150] are the main disciplines that perform this evaluation. For a time, much research about cognitive abilities in animals was performed on apes. The term 'chimpocentric' was introduced as a criticism about tests that had gone from being anthropocentric to being chimpocentric. Nonetheless, in the past decades, the perspective is much more general and any species may be a subject of study for comparative psychology: mammals (apes, cetaceans, dogs and mice), birds and some cephalopods. The evaluation focusses on "basic processes", such as perception, attention, memory, associative learning and the discrimination of concepts, and recently on more sophisticated instrumental or social abilities [150].

One of the most distinctive features of animal evaluation is the use of rewards, as instructions cannot be used. This setting is very similar to the way reinforcement learning works. Animal evaluation has also brought attention to the relevance of the interface. Clearly, the same test may require very different interfaces for a dolphin and a bonobo.

Human evaluation and animal evaluation have become more integrated in the past years, and testing procedures half way between psychometrics and comparative cognition are becoming more usual. For instance, several kinds of skills are evaluated in human children and apes in [81]. In recent years, many abilities that were considered exclusively human have been found to some extent in many animals.

Does the enlargement from humans to the whole animal kingdom suggest that these tests for animals can be used for machines? While the lower ranges of the studied abilities and the use of rewards can facilitate its application to AI systems significantly, we still have many issues about whether they can be applied to machines (at least directly). First, the selection of tasks and abilities is not systematic. Second, many of the tasks that are applied to animals would be too easy for machines (e.g., memory). And third, others would be too difficult (e.g., orientation, recognition and interaction in the real world). Nonetheless, there is an increasing need for the evaluation of animats [185] and the evaluation procedures for animals are the first candidates to try.

3.3 Evaluation using AIT

A radically different approach to AI evaluation started in the late 1990s. If intelligence was viewed as a "kind of information processing" [26] then it seemed reasonable to look at information theory for an "essential nature or formal basis of intelligence and the proper theoretical framework for it" [26]. This was finally done with *algorithmic information theory* (AIT), and the related notions of Solomonoff universal probability [155], Kolmogorov complexity [113] and Wallace's Minimum Message Length (MML) [174, 175].

There are several good properties about *algorithmic information theory* for evaluation. First, several definitions of information and complexity can be defined exclusively in computational terms, actually relative to a Universal Turing Machine (UTM), a fundamental and universal model of effective computation. For instance, the Kolmogorov complexity of an object (expressed as a binary string) relative to a UTM is defined as the shortest program (for that machine) that describes/outputs the object. Even if these definitions depend on the UTM that is used, the invariance theorem states that their values will only differ with respect to other UTM up to a constant that only depends on the two different UTMs (because one can emulate the other) [113]. The notion of algorithmic probability, introduced by Solomonoff, allows a universal distribution to be defined for each UTM, which is just the probability of objects as outputs of a UTM fed by a fair coin. While, in general, this means that compressible strings are more likely than incompressible ones, it can be shown that every computable probability distribution can be approximated by a universal distribution. In a way, Solomonoff, the father of algorithmic probability [155] gave a theoretical backing to Occam's razor. There are reasons to think that many phenomena and, as a result, many of the problems that we face every day, follow a universal distribution. This is directly linked to equation 1 again, and the discussion about the choice of the probability p. Also, we have the relevant fact, which is very significant for evaluation as well, that all universal distributions are immune to the no-free-lunch theorems, where system performance can differ very significantly for induction [107, 82]. And finally, Kolmogorov complexity and algorithmic probability are two sides of the same coin, which led to a formal connection of compression and inductive inference. It has been acknowledged that Solomonoff "solved the problem of induction" [154, 38]. Of course, not everything in AIT is straightforward. For instance, some of these concepts lead to incomputable functions, although approximations exist, such as Levin's Kt [111].

The application of AIT to (artificial) intelligence evaluation started with a variant of the Turing Test that featured compression problems [39, 40] to make the test more sufficient. While one of the goals of this work was to criticise Searle's Chinese room [147] (an argument that has faded with time), this is one of the first intelligence test proposals using AIT. At roughly the same time, a formal definition of intelligence in the form of a so-called C-test was derived from AIT [80, 65]. Figure 6 shows examples of sequences that appear in this test. They clearly resemble some exercises found in IQ tests. The major differences are that (1) sequences are obtained by a generator (a UTM with some post-conditions about the generated sequence, ensuring the unquestionability of the series continuation and less dependency on the reference machine) and (2) the fact that each sequence is accompanied by a theoretical assessment of difficulty (a variant of Levin's Kt complexity). Note the implications for evaluation of such a test, as exercises are derived from first principles (instead of being contrived by psychometricians) and the difficulty of these exercises is intrinsic, and not based on how difficult humans find them. Finally, these sequences were used to define a test by aggregating results in a way that highly resembles our recurrent equation 1, where M is formally defined as

$k = 9$	:	a, d, g, j, ...	Answer: m
$k = 12$	:	a, a, z, c, y, e, x, ...	Answer: g
$k = 14$	:	c, a, b, d, b, c, c, e, c, d, ...	Answer: d

Figure 6: Several series of different complexity 9, 12, and 14 used in the C-test [65].

including *all* possible sequences (following some conditions) and the probability is defined to cover a range of difficulties, leading to a difficulty-driven sampling as in Figure 1 (right).

Some preliminary experimental results showed that human performance correlated with the absolute difficulty (k) of each exercise and also with IQ test results for the same subjects. This encourages the use of this approach for IQ-test re-engineering. With the aim of a more complete test for machines, some extensions of the C-test were suggested, such as transforming it to work with interactive agents ("cognitive agents [...] with input/output devices for a complex environment" [80] where "rewards and penalties could be used instead" [66]) or extending them for other cognitive abilities [67]. Despite its explanatory power about IQ tests, this line of research was sharply dashed in 2003 (at least as general intelligence tests for machines) by the evidence that very simple —non-intelligent— programs could pass IQ tests [140], as we have discussed in section 3.2 (see Table 4).

Nonetheless, the extension to interactive agents was performed anyway. Interestingly, when agents and environments are considered in terms of equation 1, we just find a performance aggregation over a set of environments, exactly as had been formulated several times in the past: "intelligence is the ability of a decision-making entity to achieve success in a variety of goals when faced with a range of environments" [55]. Note that this roughly corresponds to the psychometric view of general intelligence as key to performance in a range (or all) cognitive tasks. A crucial aspect was then to define this *range* of environments, i.e., the choice of the distribution in equation 1. One option was to include *all* environments. In order to do this in a meaningful, elegant way (and get rid of any no-free lunch theorem), AIT and reinforcement learning were combined [109]. Equation 1 was instantiated with all environments as tasks with a universal distribution for p, i.e., $p(\mu) = 2^{-K(\mu)}$, with $K(\mu)$ being the Kolmogorov complexity of each environment μ. Another approach was to include all environments up to a given size or complexity, and a limit of steps [36, 37].

These proposals present several problems. First, some constructions are not computable, so approximations need to be used. Second, most environments are not really discriminative, and all agents will score the same, will just 'die' or be stuck after a few steps (this issue is partially addressed, with the use of ergodic environments [109] or world(s) where agents cannot make fatal mistakes [36]). Third, overweighting very small environments (by the use of a universal distribution or a complexity limit) makes the definition very dependent on the reference machine chosen as environment generator. Finally, time (or speed) is not considered for the environment or for the agent. For more details about these (and other) issues and some possible solutions, the reader is referred to [82] and [71, secs. 3.3 and 4]. Taking into account these solutions, some actual tests have been developed [92, 91, 94, 110]. While the results may still be useful to rank some state-of-the-art machines, if they are not compared to humans (or animals), as we discuss in the following section, the validation (or more precisely the refutation) of these tests as true intelligence tests cannot be done.

Summing up, the AIT approach is characterised by the definition of tests from formal information-based principles. This is in stark contrast to other approaches where tasks are collected, refined by trial-and-error or invented in a more arbitrary way. Most of the approaches to AI evaluation using AIT seen above have aimed at defining and measuring *general* intelligence, which is placed at the very top of the hierarchy of abilities (and hence at the opposite extreme from a specialised task-oriented evaluation). However, many interesting things can happen if AIT is applied at other layers of the hierarchy, for general cognitive abilities other than intelligence, as suggested in [67] for the passive case and hinted in [71, secs. 6.5 and 7.2] for the dynamic cases, with the use of different kinds of videogames as environments (two of the most recently introduced benchmarks and competitions are in this direction [12, 143]). Finally, the information-theoretic approach is not isolated from some of the approaches seen so far in section 2. Actually, some hybridisations and integrated approaches have been proposed [74, 43, 78, 77, 90] (apart from the compression-enriched Turing Tests ([39, 40] already mentioned above).

3.4 Universal psychometrics

Figure 7 shows the fragmentation of the approaches seen in previous sections. As we see, this fragmentation is originated by the kind of measurement we are interested in (task-oriented or ability-oriented, collected or AIT-derived tests) but most especially by the kind of subject that is being measured. In [71], the notion of 'universal test' is introduced, as a test that is applicable to "any biological or artificial system that exists at this time or in the future": human, non-human animal, enhanced human, machine, hybrid or collective. The stakes were set high, as the tests should work without knowledge about the subject, derive from computational principles, be unbiased (species, culture, language, ...), require no human intervention, be practical, produce a meaning score, and be anytime (the more time we have for the test the higher the reliability of the score). Note that in order to apply the same test to several subjects we are allowed to customise the interface, provided the features and difficulty of the items are remained unaltered. Also, we need to think about the speed of the subject, and adapt to it accordingly. Also, the capabilities of the subject can be quite varied, so the ranges of difficulty need to adapt to the agent. That suggests that universal tests must necessarily be adaptive.

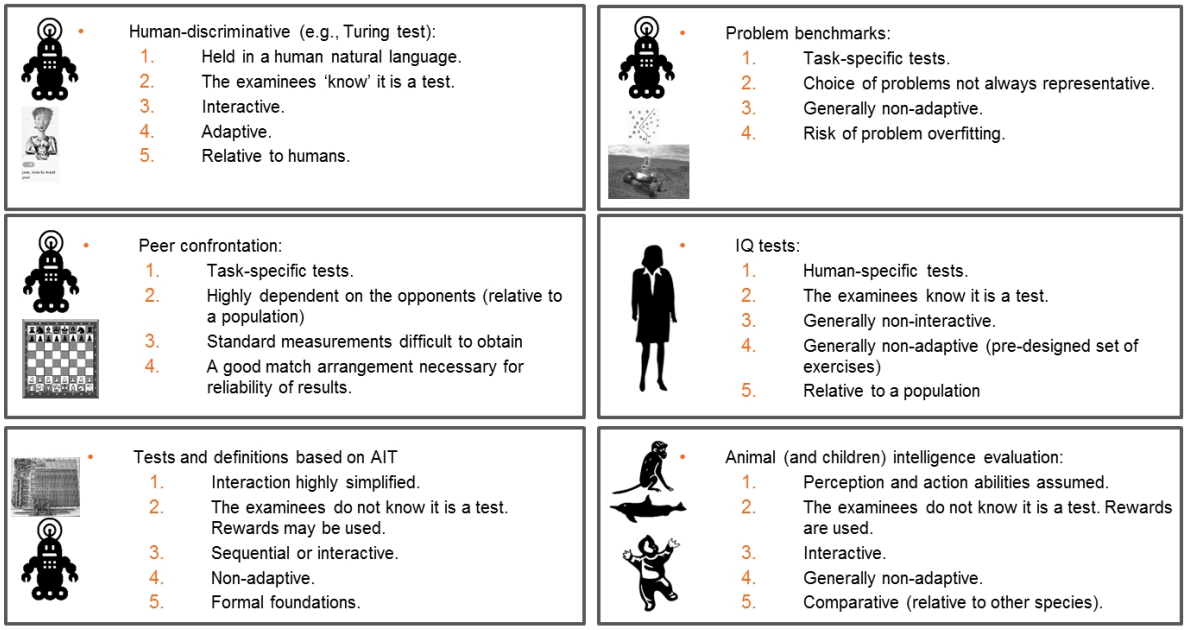
Human-discriminative (e.g., Turing test):
1. Held in a human natural language.
2. The examinees 'know' it is a test.
3. Interactive.
4. Adaptive.
5. Relative to humans.
Problem benchmarks:
1. Task-specific tests.
2. Choice of problems not always representative.
3. Generally non-adaptive.
4. Risk of problem overfitting.
Peer confrontation:
1. Task-specific tests.
2. Highly dependent on the opponents (relative to a population)
3. Standard measurements difficult to obtain
4. A good match arrangement necessary for reliability of results.
IQ tests:
1. Human-specific tests.
2. The examinees know it is a test.
3. Generally non-interactive.
4. Generally non-adaptive (pre-designed set of exercises)
5. Relative to a population
Tests and definitions based on AIT
1. Interaction highly simplified.
2. The examinees do not know it is a test. Rewards may be used.
3. Sequential or interactive.
4. Non-adaptive.
5. Formal foundations.
Animal (and children) intelligence evaluation:
1. Perception and action abilities assumed.
2. The examinees do not know it is a test. Rewards are used.
3. Interactive.
4. Generally non-adaptive.
5. Comparative (relative to other species).

Figure 7: A schematic representation of the fragmentation of the different approaches for intellgience evaluation depending of the kind of intelligent systems.

A first framework for universal, anytime intelligence tests is introduced in [71], where a class of environment is carefully chosen to be discriminative. The test starts with very simple environments and adapts to the subject's performance and speed. In this regard, this resembles a difficulty-driven sampling as described in section 2.3. The set of tasks (environments) were developed upon some of the ideas about using AIT for intelligence evaluation, as seen in section 3.3. Some experiments were performed [92, 91, 94], using the environment class defined in [68]. Difficulty was estimated using a variant of Levin's Kt. As a way of checking whether the results were meaningful, the same test compared Q-learning [178] with humans. Two different interfaces were designed on purpose. The test gave consistent results for Q-learning and humans when considered separately, but were less reasonable when put together. The experimental settings featured many limitations (simplifications, non-adaptiveness, absence of noise, low-complexity patterns, no incrementality, no social behaviour, etc.) and, probably because of this, the results did not show the actual difference between Q-learning and humans. Despite the limited results, the experiment had quite a repercussion [101, 15, 73, 191]. Nonetheless, the tests were a first effort towards a universal test and highlighted some of the challenges.

One concern about a generator of environments is the lack of richness of interaction and social behaviours

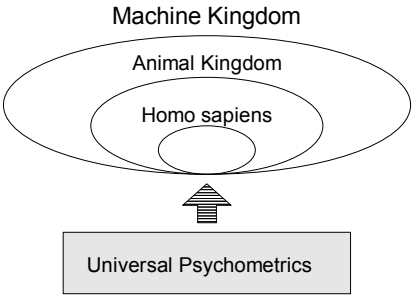

Figure 8: The realm of evaluable subjects for universal psychometrics.

that is expected. In other words, an environment that is randomly generated will have an extremely low probability of showing some social behaviour, which is a distinctive trait of human intelligence. This has suggested other ways of generating the environments and ways of incorporating other agents into them (e.g., the Darwin-Wallace distribution [74]), but it is still an open research question how to adapt these ideas to the measurement of social intelligence and multi-agent systems [43, 90, 70, 93].

The fragmentation of Figure 7 and the need of solving many of the above issues has suggested the introduction of a new perspective, dubbed 'universal psychometrics' [75]. Universal psychometrics focusses on the measurement of cognitive abilities for the 'machine kingdom', which comprises any (cognitive) system, individual or collective, either artificial, biological or hybrid. This comprehensive view is born with many hurdles ahead. Evaluation is always harder the less we know about the subject. The less we take for granted about the subjects the more difficult it is to construct a test for them. For instance, human intelligence evaluation (psychometrics) works because it is highly specialised for humans. Similarly, animal testing works (relatively well) because tests are designed in a very specific way to each species. And some of the AI evaluation settings we have already seen work because they are specialised for some kind of AI systems that are designed for some specific applications. In the case of AI, who would try to tackle a more general problem (evaluating any system) instead of the actual problem (evaluating machines)? The answer to this question is that the *actual* problem for AI is the *universal* problem. Notions such as 'animat' [185], machine-enhanced humans [30], human-enhanced machines [171], other kind of hybrids and, most especially, collectives [136] of any of the former, suggest that the distinction between animals, humans and machines is not only inappropriate, but no longer useful to advance in the evaluation of cognitive abilities. The notion of 'machine kingdom', as illustrated in Figure 8, is not very surprising to the current scientific paradigm but clarifies which class of subjects is most comprehensive.

Universal psychometrics attempts to integrate and standardise a series of concepts. A subject is seen as a physically computable (resource-bounded) interactive system. Cognitive tasks are seen as physically computable interactive systems with a score function. Interfaces are defined between subjects and tasks (observations-outputs, actions-inputs). Cognitive abilities are seen as properties over set of cognitive tasks (or task classes). As a result, the separation between task-specific and ability-specific becomes a progressive thing, depending on the generality of the class. Distributions are defined over task classes and results as aggregated performance on a task class (again, a generalised version of equation 1). Difficulty functions are computationally defined from each task. Overall, some of these elements found in psychometrics, comparative cognition and AI evaluation are overhauled here with the theory of computation and AIT. As a result, cognitive abilities are no longer *what the cognitive tests measure*, as in human psychometrics (adapting the (in)famous statement that intelligence is "what intelligence test measures" [17]), but they are properties that emanate from (general) classes of tasks, perfectly defined in computational terms. As a consequence, the relation between abilities can be explored experimentally, but also theoretically, and measures are absolute and not relativised wrt. a population (except for social abilities). This could imply some revitalisation of the white-box approach, especially for those AI systems that can be formally described in a theoretical way (e.g., some results in [88] and [82] take a white-box evaluation approach).

This view of a cognitive ability is consistent with its association with a "class of cognitive tasks" [25] that must be 'representative' for the ability. From the association between abilities and classes of tasks, "we see that by merging two cognitive task classes we get a more general cognitive task class, and a more general ability. Typically, this is studied in a hierarchical way, starting with the so-called elementary cognitive tasks

[25, page 11] (closely related to the notion of primary mental abilities of [163])" [75]. This redraws our dilemma between task-oriented and ability-oriented into a gradual hierarchy from specific tasks to general abilities, with general intelligence at the very top (including *all* possible cognitive tasks, i.e., *all* interactive Turing machines with a score function). The questions about how to sample from a task class for an effective evaluation can be generalised from our discussion in section 2.

This sets a dual view of cognitive tasks on one hand and cognitive systems on the other hand, where both spaces (the ability space and the machine kingdom) can be explored. Interestingly, both cognitive tasks and cognitive systems are defined as interactive systems, reflecting a duality world-agent. One singularity of cognitive systems (as well as their environments they are in) is that they can evolve with time, and their abilities can change. In other words, it seems that some abilities need to be constructed on top of other previously consolidated abilities, and this seems to be independent of the subject to some extent, in the same way that it seems difficult to be able to multiply without being able to add. A theoretical analysis of ability interdependency, how they can develop and the notion of potential intelligence, are still in a very incipient stage [86, 87, 72].

There can be objections and disagreements about the way many of the above concepts should be understood and defined. There can also be objections about what a universal test should look like [42]. But a more integrated view of cognitive abilities for humans, animals, robots, agents, animats, hybrids, swarms, etc., is not only possible but useful. Bear in mind that universal psychometrics does not exclude the use of non-universal tests, as tests that are non-universal can be more efficient (tests can be universal or not, depending on the application), but aims at having a more integrated and well-founded view of how intelligent systems are evaluated in terms of cognitive abilities.

4 Conclusions

We started this paper looking at the way AI evaluation is commonly performed, through task-oriented evaluation, mostly with a black-box approach. We identified several problems and limitations, and we noticed that there is still a huge margin of improvement in the way AI systems are evaluated. The key issues are the set of tasks M and their distribution p, as well as distinguishing the definition of the problem class (aggregation) from an effective sampling procedure (testing procedure). Then we switched to ability-oriented evaluation, a much more immature approach, but that may have a more relevant role in the future. The notion and evaluation of abilities is more elusive than the notion and evaluation of tasks. We have argued that this requires the integration of several perspectives that are currently scattered efforts in AI, psychometrics, AIT and comparative cognition. The different areas, philosophies, tools, foundations, terminologies and the different kinds of subjects to be evaluated can be unified with an integrated perspective known as universal psychometrics. Here, the exploration of the machine kingdom is dual to the exploration of the set of possible cognitive abilities/tasks. In both spaces we aim at becoming more general, which is where evaluation is more challenging (see Figure 9). This resembles the duality in the theory of computation (e.g., problem classes and automata classes). The more formal approach advocated by universal psychometrics can make the white-box evaluation approach recover some relevance in AI.

From the problems and limitations found in AI evaluation and the tools and ideas that have appeared along the paper, we now enumerate a number of generic guidelines. These can be considered when an AI evaluation setting is under consideration.

- The definition of Ω, the set of possible systems that can be evaluated (or that can be opponents in peer confrontation evaluation), must be clarified from the beginning. Any information about their proficiency and expected characteristics may be very useful. If humans are considered, the way in which they are admitted and how they are instructed must be defined. The more general Ω is the less we can assume about the evaluation process. If Ω is heterogeneous (e.g., a universal test), different interfaces must be considered.

- The definition of M, the set of possible tasks, and its associated distribution p configure what we are measuring. This can be built from a set of problems or using a generator. This pair $\langle M, p \rangle$ has to be representative of a task (in task-oriented evaluation) or an ability (in ability-oriented evaluation). If

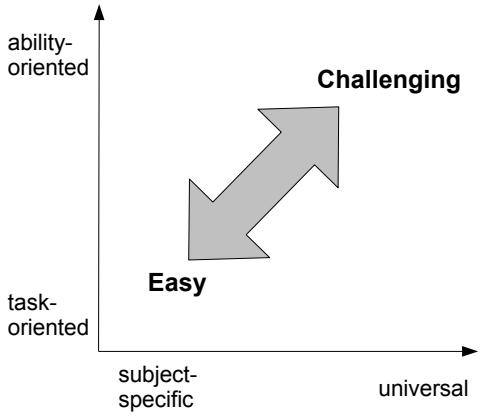

Figure 9: Tests become more or less challenging depending on the generality of the class of subjects under consideration (from subject-specific to universal) and the class of abilities (from task-oriented evaluation to ability-oriented evaluation).

it is a peer confrontation evaluation, M will be enlarged with as many combinations between game (environment) and agents in Ω are possible. The distribution p will be updated accordingly.

- The definition of R and its aggregation Φ must ensure that the values $R(\mu)$ for all $\mu \in M$ are going to be commensurate and that the aggregation is bounded. An analysis about expected measurement error is useful at this point. The robustness of R depending on the length or time left for each episode will indicate whether repetitions are needed to reduce the measurement error given by $R(\mu)$.

- As much as possible, the similarity between tasks or a set of features describing them should be identified. An intrinsic difficulty function (even if approximate) is always very useful. Showing the distribution of difficulty for M can be highly informative. If difficulty is available, item response curves could be prepared.

- The sampling method must be as much efficient as possible, by using, e.g., a clustering sampling or a range of difficulties if we have a non-adaptive evaluation. For the peer-confrontation evaluation, the arrangement of matches can be designed beforehand if the evaluation is not adaptive. Similarly, the procedure for an adaptive evaluation must also be carefully designed to ensure measurement robustness. Simulations can be useful to estimate this.

- Information about how the evaluation is performed (including R, Φ and some illustrative problems) can be disclosed to the systems that are being evaluated (or to their designers). However, Ω, M and p should not be disclosed. If possible, the problems should not be disclosed after the evaluation either, as keeping them secret makes it possible to compare with the same problems for different subjects or at different times (e.g., we can evaluate progress of a system or a discipline during a period).

- After the evaluation, results must be analysed beyond the mere calculation of the aggregated results. Item response functions and agent response functions [75] can be constructed empirically from the results and compared with the theoretical functions or any other information about Ω and M. Discrepancies or anomalies may suggest that the evaluation setting has to be revised. Results of the evaluation must become public at the highest possible detail, so they can be analysed and compared by other researchers and participants (following, e.g., the notion of 'experiment database' [168], such as in the machine learning community[51]).

It is of course an open question to what extent the above recommendations will be followed on a regular basis for AI evaluation. It can be argued whether AI evaluation has been a priority for AI in the past, but it seems that it has not been recognised as an imperative problem or a mainstream area of research. If this is the case, this paper can help change this perspective. Anyhow, the question of AI evaluation remains and there

[51] http://openml.org/.

is space for significant improvement, even for the most specific sets Ω and M (bottom-left part of Figure 9). At the other end, measuring intelligence and doing it universally is a key ingredient for understanding what intelligence is (and, of course, to devise intelligent artefacts). Many interesting questions and applications lay in the middle of Figure 9, as AI evaluation is no longer limited to task-specific evaluation of AI systems or to evaluating progress in AI. Instead, AI is becoming able to evaluate systems that learn to solve instead of systems that are programmed to solve.

In any case, and with any of the approaches seen so far, a more scientific theory of AI evaluation is being required for many applications (CAPTCHAs, social networks, agent certification, etc.) and it will be more and more common in a future with a plethora of bots, robots, artificial agents, avatars, control systems, 'animats', hybrids, collectives, etc. It is also crucial for the technological singularity once (and if) achieved [45], especially because some of the prophecies and forecasts disregard that the first thing to consider about the singularity is to have metrics to detect whether and where AI progresses towards it.

Summing up, AI requires an accurate, effective, non-anthropocentric, meaningful and computational way of evaluating its progress, by evaluating its artefacts. This paper can serve as a comprehensive source of the state of the art of the AI evaluation, its challenges and the avenues for future work.

Acknowledgements

This work was supported by the EU (FEDER) and the Spanish MINECO projects CONSOLIDER-INGENIO CSD2007-00022, TIN 2010-21062-C02-02, and TIN 2013-45732-C4-1-P, by Generalitat Valenciana projects Prometeo/2008/051 and PROMETEO/2011/052, and the REFRAME project granted by the European Coordinated Research on Long-term Challenges in Information and Communication Sciences & Technologies ERA-Net (CHIST-ERA), and funded by Ministerio de Economía y Competitividad, with code PCIN-2013-037. I thank the organisers of the Summer School of the Spanish Association for Artificial Intelligence, in A Coruña, Spain, held in September 2014, for giving me the opportunity to give a lecture on 'AI Evaluation'. This paper evolved in parallel with that lecture. The coverage of the BotPrize competition was discussed with Manuel González-Bedia. Figure 5 is courtesy of Fernando Martínez-Plumed.

References

[1] J. Alcalá, A. Fernández, J. Luengo, J. Derrac, S. García, L. Sánchez, and F. Herrera. Keel data-mining software tool: Data set repository, integration of algorithms and experimental analysis framework. *Journal of Multiple-Valued Logic and Soft Computing*, 17:255–287, 2010. 11

[2] J. R. M. Alexander and S. Smales. Intelligence, learning and long-term memory. *Personality and Individual Differences*, 23(5):815–825, 1997. 17

[3] T. Alpcan, T. Everitt, and M. Hutter. Can we measure the difficulty of an optimization problem? *IEEE Information Theory Workshop (ITW)*, 2014. 5

[4] R. Alur, R. Bodik, G. Juniwal, M. M. K. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *Formal Methods in Computer-Aided Design (FMCAD), 2013*, pages 1–17. IEEE, 2013. 11

[5] N. Alvarado, S. S. Adams, S. Burbeck, and C. Latta. Beyond the Turing Test: Performance metrics for evaluating a computer simulation of the human mind. In *Development and Learning, 2002. Proceedings. The 2nd International Conference on*, pages 147–152. IEEE, 2002. 6

[6] J. Anderson, J. Baltes, and C. T. Cheng. Robotics competitions as benchmarks for AI research. *The Knowledge Engineering Review*, 26(01):11–17, 2011. 2, 11

[7] I. Arel, D. C. Rose, and T. P. Karnowski. Deep machine learning - a new frontier in artificial intelligence research. *Computational Intelligence Magazine, IEEE*, 5(4):13–18, 2010. 2

[8] M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida. Cognitive developmental robotics: a survey. *Autonomous Mental Development, IEEE Transactions on*, 1(1):12–34, 2009. 2

[9] K. Bache and M. Lichman. UCI machine learning repository, 2013. 7, 11, 12

[10] A. J. Bagnall and Z. V. Zatuchna. On the classification of maze problems. In *Foundations of Learning Classifier Systems*, pages 305–316. Springer, 2005. 8

[11] D. Baldwin and S. B. Yadav. The process of research investigations in artificial intelligence - a unified view. *Systems, Man and Cybernetics, IEEE Transactions on*, 25(5):852–861, 1995. 2

[12] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 06 2013. 7, 11, 21

[13] J. L. Benacloch-Ayuso. Integration of general game playing with RL-glue. Technical report, DSIC, Universitat Politècnica de València, 2012. 15

[14] T. R. Besold. A note on chances and limitations of psychometric ai. In *KI 2014: Advances in Artificial Intelligence*, pages 49–54. Springer, 2014. 18

[15] C. Biever. Ultimate IQ: one test to rule them all. *New Scientist*, 211(2829,10 September 2011):42–45, 2011. 22

[16] M. Borg, S. S. Johansen, D. L. Thomsen, and M. Kraus. Practical implementation of a graphics Turing Test. In *Advances in Visual Computing*, pages 305–313. Springer, 2012. 6

[17] E. G. Boring. Intelligence as the tests test it. *New Republic*, pages 35–37, 1923. 23

[18] N. Bostrom. *Superintelligence: Paths, dangers, strategies*. Oxford University Press, 2014. 16

[19] S. Bringsjord. Psychometric artificial intelligence. *Journal of Experimental & Theoretical Artificial Intelligence*, 23(3):271–277, 2011. 18

[20] S. Bringsjord and B. Schimanski. What is artificial intelligence? Psychometric AI as an answer. In *International Joint Conference on Artificial Intelligence*, pages 887–893, 2003. 18

[21] B. G. Buchanan. *Artificial intelligence as an experimental science*. Springer, 1988. 2, 5

[22] M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon's mechanical turk a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011. 19

[23] E. Bursztein, J. Aigrain, A. Moscicki, and J. C. Mitchell. The end is nigh: generic solving of text-based captchas. In *Proceedings of the 8th USENIX conference on Offensive Technologies*, pages 3–3. USENIX Association, 2014. 6

[24] M. Campbell, A. J. Hoane, and F. Hsu. Deep Blue. *Artificial Intelligence*, 134(1-2):57 – 83, 2002. 3

[25] J. B. Carroll. *Human cognitive abilities: A survey of factor-analytic studies*. Cambridge University Press, 1993. 23, 24

[26] B Chandrasekaran. What kind of information processing is intelligence? In *The foundation of artificial intelligence—a sourcebook*, pages 14–46. Cambridge University Press, 1990. 20

[27] Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia. Who is tweeting on twitter: human, bot, or cyborg? In *Proceedings of the 26th annual computer security applications conference*, pages 21–30. ACM, 2010. 6

[28] W. G. Cochran. *Sampling techniques*. John Wiley & Sons, 2007. 8

[29] P. R. Cohen and A. E. Howe. How evaluation guides AI research: The message still counts more than the medium. *AI Magazine*, 9(4):35, 1988. 2, 3

[30] Y. Cohen. Testing and cognitive enhancement. Technical report, National Institute for Testing and Evaluation, Jerusalem, Israel, 2013. 23

[31] J. G. Conrad and J. Zeleznikow. The significance of evaluation in AI and law: a case study re-examining ICAIL proceedings. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pages 186–191. ACM, 2013. 2

[32] I. J. Deary, G. Der, and G. Ford. Reaction times and intelligence differences: A population-based cohort study. *Intelligence*, 29(5):389–399, 2001. 19

[33] K. S. Decker, E. H. Durfee, and V. R. Lesser. Evaluating research in cooperative distributed problem solving. *Distributed Artificial Intelligence*, 2:487–519, 1989. 2

[34] J. Demšar. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research*, 7:1–30, 2006. 12

[35] D. K. Detterman. A challenge to Watson. *Intelligence*, 39(2-3):77 – 78, 2011. 18

[36] D. Dobrev. AI - What is this? A definition of artificial intelligence. *PC Magazine Bulgaria (in Bulgarian, English version at http://www.dobrev.com/AI)*, 2000. 21

[37] D. Dobrev. Formal definition of artificial intelligence. *International Journal of Information Theories and Applications*, 12(3):277–285, 2005. 21

[38] D. L Dowe. Introduction to Ray Solomonoff 85th memorial conference. In D. L. Dowe, editor, *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, volume 7070 of *Lecture Notes in Computer Science*, pages 1–36. Springer Berlin Heidelberg, 2013. 20

[39] D. L. Dowe and A. R. Hajek. A computational extension to the Turing Test. In *Proceedings of the 4th Conference of the Australasian Cognitive Science Society, University of Newcastle, NSW, Australia*, 1997. 20, 21

[40] D. L. Dowe and A. R. Hajek. A non-behavioural, computational extension to the Turing Test. In *Intl. Conf. on Computational Intelligence & multimedia applications (ICCIMA'98), Gippsland, Australia*, pages 101–106, 1998. 20, 21

[41] D. L. Dowe and J. Hernández-Orallo. IQ tests are not for machines, yet. *Intelligence*, 40(2):77–81, 2012. 19

[42] D. L. Dowe and J. Hernández-Orallo. How universal can an intelligence test be? *Adaptive Behavior*, 22(1):51–69, 2014. 24

[43] D. L. Dowe, J. Hernández-Orallo, and P. K. Das. Compression and intelligence: social environments and communication. In J. Schmidhuber, K.R. Thórisson, and M. Looks, editors, *Artificial General Intelligence*, volume 6830, pages 204–211. LNAI series, Springer, 2011. 21, 23

[44] C. Drummond and N. Japkowicz. Warning: statistical benchmarking is addictive. Kicking the habit in machine learning. *Journal of Experimental & Theoretical Artificial Intelligence*, 22(1):67–80, 2010. 2, 7, 12

[45] A. H. Eden, J. H. Moor, J. H. Soraker, and E. Steinhart. *Singularity hypotheses: A scientific and philosophical assessment*. Springer, 2013. 26

[46] A. E. Elo. *The rating of chessplayers, past and present*, volume 3. Batsford London, 1978. 5, 13

[47] S. E. Embretson and S. P. Reise. *Item response theory for psychologists*. L. Erlbaum, 2000. 8, 9, 18

[48] J. M. Evans and E. R. Messina. Performance metrics for intelligent systems. *NIST Special Publication SP*, pages 101–104, 2001. 15

[49] T. Everitt, T. Lattimore, and M. Hutter. Free lunch for optimisation under the universal distribution. In *Evolutionary Computation (CEC), 2014 IEEE Congress on*, pages 167–174. IEEE, 2014. 4

[50] E. Falkenauer. On method overfitting. *Journal of Heuristics*, 4(3):281–287, 1998. 2, 7

[51] P. J. Ferrando. Difficulty, discrimination, and information indices in the linear factor analysis model for continuous item responses. *Applied Psychological Measurement*, 33(1):9–24, 2009. 9

[52] P. J. Ferrando. Assessing the discriminating power of item and test scores in the linear factor-analysis model. *Psicológica*, 33:111–139, 2012. 9

[53] C. Ferri, J. Hernández-Orallo, and R Modroiu. An experimental comparison of performance measures for classification. *Pattern Recognition Letters*, 30(1):27–38, 2009. 12

[54] D. Ferrucci, E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. A. Kalyanpur, A. Lally, J.W. Murdock, E. Nyberg, J. Prager, et al. Building Watson: An overview of the DeepQA project. *AI Magazine*, 31(3):59–79, 2010. 18

[55] D. B. Fogel. The evolution of intelligent decision making in gaming. *Cybernetics and Systems*, 22(2):223–236, 1991. 21

[56] J. Gaschnig, P. Klahr, H. Pople, E. Shortliffe, and A. Terry. Evaluation of expert systems: Issues and case studies. *Building expert systems*, 1:241–278, 1983. 2

[57] J. R. Geissman and R. D. Schultz. Verification & validation. *AI Expert*, 3(2):26–33, 1988. 2

[58] M. Genesereth, N. Love, and B. Pell. General game playing: Overview of the AAAI competition. *AI Magazine*, 26(2):62, 2005. 14

[59] D. Gerónimo and A. M. López. Datasets and benchmarking. In *Vision-based Pedestrian Protection Systems for Intelligent Vehicles*, pages 87–93. Springer, 2014. 11

[60] B. Goertzel and C. Pennachin, editors. *Artificial general intelligence*. Springer, 2007. 2

[61] O. Goldreich and S. Vadhan. Special issue on worst-case versus average-case complexity – editors' foreword. *computational complexity*, 16(4):325–330, 2007. 4

[62] S. Gulwani, J. Hernández-Orallo, E. Kitzelmann, S. H. Muggleton, U. Schmid, and B. Zorn. Inductive programming meets the real world. *Submitted*, 2014. 2

[63] S. Gulwani, E. Kitzelmann, and U. Schmid. Approaches and applications of inductive programming (dagstuhl seminar 13502). *Dagstuhl Reports*, 3(12), 2014. 2

[64] D. J. Hand. *Measurement theory and practice*. A Hodder Arnold Publication, 2004. 2

[65] J. Hernández-Orallo. Beyond the Turing Test. *J. Logic, Language & Information*, 9(4):447–466, 2000. 20, 21

[66] J. Hernández-Orallo. On the computational measurement of intelligence factors. In A. Meystel, editor, *Performance metrics for intelligent systems workshop*, pages 1–8. National Institute of Standards and Technology, Gaithersburg, MD, U.S.A., 2000. 5, 21

[67] J. Hernández-Orallo. Thesis: Computational measures of information gain and reinforcement in inference processes. *AI Communications*, 13(1):49–50, 2000. 21

[68] J. Hernández-Orallo. A (hopefully) non-biased universal environment class for measuring intelligence of biological and artificial systems. In M. Hutter et al., editor, *Artificial General Intelligence, 3rd Intl Conf*, pages 182–183. Atlantis Press, Extended report at http://users.dsic.upv.es/proy/anynt/unbiased.pdf, 2010. 22

[69] J. Hernández-Orallo. Deep knowledge: Inductive programming as an answer. *Approaches and Applications of Inductive Programming (Dagstuhl Seminar 13502), Gulwani, S. and Kitzelmann, E. and Schmid, U. (eds.)*, 2014. 2

[70] J. Hernández-Orallo. On environment difficulty and discriminating power. *Autonomous Agents and Multi-Agent Systems*, pages 1–53, 2014. 13, 14, 23

[71] J. Hernández-Orallo and D. L. Dowe. Measuring universal intelligence: Towards an anytime intelligence test. *Artificial Intelligence*, 174(18):1508 – 1539, 2010. 21, 22

[72] J. Hernández-Orallo and D. L. Dowe. On potential cognitive abilities in the machine kingdom. *Minds and Machines*, 23:179–210, 2013. 24

[73] J. Hernández-Orallo and D. L. Dowe. Mammals, machines and mind games. Who's the smartest? *The Conversation*, http://theconversation.edu.au/articles/mammals-machines-and-mind-games-whos-the-smartest-1125, April 2011. 22

[74] J. Hernández-Orallo, D. L. Dowe, S. España-Cubillo, M. V. Hernández-Lloreda, and J. Insa-Cabrera. On more realistic environment distributions for defining, evaluating and developing intelligence. In J. Schmidhuber, K.R. Thórisson, and M. Looks, editors, *Artificial General Intelligence*, volume 6830, pages 82–91. LNAI, Springer, 2011. 21, 23

[75] J. Hernández-Orallo, D. L. Dowe, and M. V. Hernández-Lloreda. Universal psychometrics: Measuring cognitive abilities in the machine kingdom. *Cognitive Systems Research*, 27:5074, 2014. 23, 24, 25

[76] J. Hernández-Orallo, P. Flach, and C. Ferri. A unified view of performance metrics: Translating threshold choice into expected classification loss. *The Journal of Machine Learning Research*, 13(1):2813–2869, 2012. 12

[77] J. Hernández-Orallo, J. Insa, D. L. Dowe, and B. Hibbard. Turing Tests with Turing machines. In Andrei Voronkov, editor, *Turing-100*, volume 10, pages 140–156. EPiC Series, 2012. 6, 21

[78] J. Hernández-Orallo, J. Insa-Cabrera, D.L. Dowe, and B. Hibbard. Turing machines and recursive Turing Tests. In V. Muller and A. Ayesh, editors, *AISB/IACAP 2012 Symposium "Revisiting Turing and his Test"*, pages 28–33. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour, 2012. 6, 21

[79] J. Hernández-Orallo, F. Martínez-Plumed, U. Schmid, M. Siebers, and D. L. Dowe. Computer models solving human intelligence test problems: progress and implications. *submitted*, 2014. 19

[80] J. Hernández-Orallo and N. Minaya-Collado. A formal definition of intelligence based on an intensional variant of Kolmogorov complexity. In *Proc. Intl Symposium of Engineering of Intelligent Systems (EIS'98)*, pages 146–163. ICSC Press, 1998. 20, 21

[81] E. Herrmann, J. Call, M. V. Hernández-Lloreda, B. Hare, and M. Tomasello. Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, Vol 317(5843):1360–1366, 2007. 20

[82] B. Hibbard. Bias and no free lunch in formal measures of intelligence. *Journal of Artificial General Intelligence*, 1(1):54–61, 2009. 20, 21, 23

[83] P. Hingston. A new design for a Turing Test for bots. In *Computational Intelligence and Games (CIG), 2010 IEEE Symposium on*, pages 345–350. IEEE, 2010. 6

[84] P. Hingston. *Believable Bots: Can Computers Play Like People?* Springer, 2012. 6

[85] T. K. Ho and M. Basu. Complexity measures of supervised classification problems. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(3):289–300, 2002. 12

[86] M. Hutter. The fastest and shortest algorithm for all well-defined problems. *International Journal of Foundations of Computer Science*, 13:431–443, 2002. 24

[87] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, 2005. 24

[88] M. Hutter. Universal algorithmic intelligence: A mathematical top→down approach. In B. Goertzel and C. Pennachin, editors, *Artificial General Intelligence*, Cognitive Technologies, pages 227–290. Springer, Berlin, 2007. 2, 23

[89] C. Igel and M. Toussaint. A no-free-lunch theorem for non-uniform distributions of target functions. *Journal of Mathematical Modelling and Algorithms*, 3(4):313–322, 2005. 4

[90] J. Insa-Cabrera, J. L. Benacloch-Ayuso, and J. Hernández-Orallo. On measuring social intelligence: Experiments on competition and cooperation. In J. Bach, B. Goertzel, and M. Iklé, editors, *AGI*, volume 7716 of *Lecture Notes in Computer Science*, pages 126–135. Springer, 2012. 21, 23

[91] J. Insa-Cabrera, D. L. Dowe, S. España-Cubillo, M. V. Hernández-Lloreda, and J. Hernández-Orallo. Comparing humans and AI agents. In J. Schmidhuber, K.R. Thórisson, and M. Looks, editors, *Artificial General Intelligence*, volume 6830, pages 122–132. LNAI, Springer, 2011. 21, 22

[92] J. Insa-Cabrera, D. L. Dowe, and J. Hernández-Orallo. Evaluating a reinforcement learning algorithm with a general intelligence test. In J.A. Moreno J.A. Lozano, J.A. Gamez, editor, *Current Topics in Artificial Intelligence. CAEPIA 2011*. LNAI Series 7023, Springer, 2011. 21, 22

[93] J. Insa-Cabrera and J. Hernández-Orallo. Definition and properties to assess multi-agent environments as social intelligence tests. *arXiv preprint, http://arxiv.org/abs/1408.6350*, 2014. 15, 23

[94] J. Insa-Cabrera, J. Hernández-Orallo, D.L. Dowe, S. Espa na, and M.V. Hernández-Lloreda. The anynt project intelligence test : Lambda - one. In V. Muller and A. Ayesh, editors, *AISB/IACAP 2012 Symposium "Revisiting Turing and his Test"*, pages 20–27. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour, 2012. 21, 22

[95] A. Jacoff, E. Messina, B. A. Weiss, S. Tadokoro, and Y. Nakagawa. Test arenas and performance metrics for urban search and rescue robots. In *Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 4, pages 3396–3403. IEEE, 2003. 11

[96] N. Japkowicz and M. Shah. *Evaluating Learning Algorithms*. Cambridge University Press, 2011. 12

[97] T. Z. Keith and M. R Reynolds. Cattell–Horn–Carroll abilities and cognitive tests: What we've learned from 20 years of research. *Psychology in the Schools*, 47(7):635–650, 2010. 17

[98] W. Ketter and A. Symeonidis. Competitive benchmarking: Lessons learned from the trading agent competition. *AI Magazine*, 33(2):103, 2012. 14

[99] J. H. Kim. *Soccer robotics*, volume 11. Springer, 2004. 14

[100] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa. Robocup: The robot world cup initiative. In *Proceedings of the first international conference on Autonomous agents*, pages 340–347. ACM, 1997. 14

[101] K. Kleiner. Who are you calling bird-brained? An attempt is being made to devise a universal intelligence test. *The Economist*, 398(8723, 5 March 2011):82, 2011. 22

[102] D. E. Knuth. *Sorting and searching, volume 3 of The Art of Computer Programming*. Addison-Wesley, 1973. 4

[103] J. R. Koza. Human-competitive results produced by genetic programming. *Genetic Programming and Evolvable Machines*, 11(3-4):251–284, 2010. 6

[104] J. Langford. Clever methods of overfitting. *Machine Learning (Theory), http://hunch.net*, 2005. 2, 7

[105] P. Langley. Research papers in machine learning. *Machine Learning*, 2(3):195–198, 1987. 2

[106] P. Langley. The changing science of machine learning. *Machine Learning*, 82(3):275–279, 2011. 2

[107] T. Lattimore and M. Hutter. No free lunch versus Occam's razor in supervised learning. In *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, pages 223–235. Springer, 2013. 4, 20

[108] S. Legg and M. Hutter. Tests of machine intelligence. In Max Lungarella, Fumiya Iida, Josh Bongard, and Rolf Pfeifer, editors, *50 Years of Artificial Intelligence*, volume 4850 of *Lecture Notes in Computer Science*, pages 232–242. Springer Berlin Heidelberg, 2007. 2

[109] S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4):391–444, 2007. 21

[110] S. Legg and J. Veness. An approximation of the universal intelligence measure. In *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, pages 236–249. Springer, 2013. 21

[111] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973. 20

[112] L. A. Levin. Average case complete problems. *SIAM J. on Computing*, 15:285–286, 1986. 4

[113] M. Li and P. Vitányi. *An introduction to Kolmogorov complexity and its applications (3rd ed.)*. Springer-Verlag, 2008. 4, 20

[114] D. Livingstone. Turing's test and believable AI in games. *Computers in Entertainment (CIE)*, 4(1):6, 2006. 6

[115] J. M. Llargues-Asensio, J. Peralta, R. Arrabales, M. González-Bedía, P. Cortez, and A. L. López-Peña. Artificial intelligence approaches for the generation and assessment of believable human-like behaviour in virtual characters. *Expert Systems with Applications*, 2014. 7

[116] D. Long and M. Fox. The 3rd international planning competition: Results and analysis. *J. Artif. Intell. Res. (JAIR)*, 20:1–59, 2003. 11

[117] F. M. Lord. *Applications of item response theory to practical testing problems*. Mahwah, NJ: Erlbaum, 1980. 18

[118] N. Macià and E. Bernadó-Mansilla. Towards UCI+: A mindful repository design. *Information Sciences*, 261:237–262, 2014. 12

[119] R. Madhavan, E. Tunstel, and E. Messina. *Performance Evaluation and Benchmarking of Intelligent Systems*. Springer, September, 2009. 2, 15

[120] C. Marché and H. Zantema. The termination competition. In *Term Rewriting and Applications*, pages 303–313. Springer, 2007. 11

[121] H. Masum and S. Christensen. The turing ratio: A framework for open-ended task metrics. In *Journal of Evolution and Technology*. Citeseer, 2003. 13, 16

[122] H. Masum, S. Christensen, and F. Oppacher. The turing ratio: Metrics for open-ended tasks. In *GECCO*, pages 973–980. Citeseer, 2002. 16

[123] J. McCarthy. What is artificial intelligence. Technical report, Stanford University, http://www-formal.stanford.edu/jmc/whatisai.html, 2007. 1

[124] P. McCorduck. *Machines who think*. A K Peters/CRC Press, 2004. 1

[125] J. McDermott, D. R. White, S. Luke, L. Manzoni, M. Castelli, L. Vanneschi, W. Jaśkowski, K. Krawiec, R. Harper, K. De Jong, and U.-M. O'Reilly. Genetic programming needs better benchmarks. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*, pages 791–798, Philadelphia, 2012. ACM. 11

[126] M. McGuigan. Graphics Turing Test. *arXiv preprint cs/0603132*, 2006. 6

[127] G. J. Mellenbergh. Generalized linear item response theory. *Psychological Bulletin*, 115(2):300, 1994. 9

[128] A. Meystel, J. Albus, E. Messina, and D. Leedom. Performance measures for intelligent systems: Measures of technology readiness. Technical report, DTIC Document, 2003. 15

[129] M. L. Minsky, editor. *Semantic Information Processing*. MIT Press, 1968. 1

[130] S. T. Mueller and B. S. Minnery. Adapting the Turing Test for embodied neurocognitive evaluation of biologically-inspired cognitive agents. In *Proc. 2008 AAAI Fall Symposium on Biologically Inspired Cognitive Architectures*, 2008. 6

[131] A. Newell. You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In *Visual Information Processing, ed. W. Chase*, pages 283–308. New York: Academic Press, 1973. 18

[132] A. Newell and H. A. Simon. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126, 1976. 2

[133] G. Oppy and D. L. Dowe. The Turing Test. In Edward N. Zalta, editor, *Stanford Encyclopedia of Philosophy*, pages Stanford University, http://plato.stanford.edu/entries/turing–test/, 2011. 5

[134] M. Potthast, M. Hagen, T. Gollub, M. Tippmann, J. Kiesel, P. Rosso, E. Stamatatos, and B. Stein. Overview of the 5th international competition on plagiarism detection. *CLEF 2013 Evaluation Labs and Workshop Working Notes Papers, 23-26 September, Valencia, Spain*, 2013. 11

[135] D. Proudfoot. Anthropomorphism and AI: Turing's much misunderstood imitation game. *Artificial Intelligence*, 175(5):950–957, 2011. 6

[136] A. J. Quinn and B. B. Bederson. Human computation: a survey and taxonomy of a growing field. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1403–1412. ACM, 2011. 23

[137] S. Rajani. Artificial intelligence – man or machine. *International Journal of Information Technology*, 4(1):173–176, 2011. 15

[138] J. Rothenberg, J. Paul, I. Kameny, J. R. Kipps, and M. Swenson. Evaluating expert system tools: A framework and methodology–workshops. Technical report, DTIC Document, 1987. 2

[139] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2009. 3, 17

[140] P. Sanghi and D. L. Dowe. A computer program capable of passing IQ tests. In *4th Intl. Conf. on Cognitive Science (ICCS'03), Sydney*, pages 570–575, 2003. 18, 19, 21

[141] J. Schaeffer, N. Burch, Y. Bjornsson, A. Kishimoto, M. Muller, R. Lake, P. Lu, and S. Sutphen. Checkers is solved. *Science*, 317(5844):1518, 2007. 4, 13

[142] K. W. Schaie. Primary mental abilities. *Corsini Encyclopedia of Psychology*, 2010. 17

[143] T. Schaul. An extensible description language for video games. *Computational Intelligence and AI in Games, IEEE Transactions on*, PP(99):1–1, 2014. 7, 11, 21

[144] C. Schenck. Intelligence tests for robots: Solving perceptual reasoning tasks with a humanoid robot. Master's thesis, Iowa State University, 2013. 19

[145] C. Schlenoff, H. Scott, and S. Balakirsky. Performance evaluation of intelligent systems at the national institute of standards and technology (nist). Technical report, DTIC Document, 2011. 2, 15

[146] P. Schweizer. The truly total Turing Test. *Minds and Machines*, 8(2):263–272, 1998. 6

[147] J. R. Searle. Minds, brains, and programs. *The Behavioral and Brain Sciences*, 3:417–457, 1980. 20

[148] G. A. F. Seber and M. M. Salehi. Adaptive cluster sampling. In *Adaptive Sampling Designs*, pages 11–26. Springer, 2013. 8

[149] S. J. Shettleworth. *Cognition, evolution, and behavior*. Oxford University Press, 2010. 19

[150] S. J. Shettleworth, P. Bloom, and L. Nadel. *Fundamentals of Comparative Cognition*. Oxford University Press, 2013. 19

[151] H. A. Simon. Artificial intelligence: an empirical science. *Artificial Intelligence*, 77(1):95–127, 1995. 2, 5

[152] W. D. Smith. Rating systems for gameplayers, and learning. *NEC, Princeton, NJ, Tech. Rep*, pages 93–104, 2002. 13

[153] C. Soares. UCI++: Improved support for algorithm selection using datasetoids. In *Advances in Knowledge Discovery and Data Mining*, pages 499–506. Springer, 2009. 12

[154] R. Solomonoff. Does algorithmic probability solve the problem of induction. *Information, Statistics and Induction in Science*, pages 7–8, 1996. 20

[155] R. J. Solomonoff. A formal theory of inductive inference. Part I. *Information and control*, 7(1):1–22, 1964. 4, 20

[156] R. Srinivasan. *Importance sampling: Applications in communications and detection*. Springer, 2002. 8

[157] B. Starkie, M. van Zaanen, and D. Estival. The Tenjinno machine translation competition. In *Grammatical Inference: Algorithms and Applications*, pages 214–226. Springer, 2006. 11

[158] R. J. Sternberg (ed.). *Handbook of intelligence*. Cambridge University Press, 2000. 18

[159] R. E. Strickler. Change in selected characteristics of students between ninth and twelfth grade as related to high school curriculum., 1973. 17

[160] N. Sturtevant. Benchmarks for grid-based pathfinding. *Transactions on Computational Intelligence and AI in Games*, 4(2):144 – 148, 2012. 8, 11

[161] G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure: The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009. 11

[162] G. Sutcliffe and C. Suttner. The State of CASC. *AI Communications*, 19(1):35–48, 2006. 11

[163] L. L. Thurstone. Primary mental abilities. *Psychometric monographs*, 1938. 24

[164] J. Togelius, G. N. Yannakakis, S. Karakovskiy, and N. Shaker. Assessing believability. In *Believable Bots*, pages 215–230. Springer, 2012. 7

[165] A. M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950. 5

[166] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984. 4

[167] J. N. van Rijn, B. Bischl, L. Torgo, B. Gao, V. Umaashankar, S. Fischer, P. Winter, B. Wiswedel, Michael R Berthold, and J. Vanschoren. Openml: a collaborative science platform. In *Machine Learning and Knowledge Discovery in Databases*, pages 645–649. Springer, 2013. 12

[168] J. Vanschoren, H. Blockeel, B. Pfahringer, and G. Holmes. Experiment databases. *Machine Learning*, 87(2):127–158, 2012. 12, 25

[169] J. Vanschoren, J. N. van Rijn, B. Bischl, and L. Torgo. Openml: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014. 12

[170] D. Vázquez, A. M. López, J. Marín, D. Ponsa, and D. Gerónimo. Virtual and real world adaptation for pedestrian detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 36(4):797–809, April 2014. 7

[171] L. von Ahn. Human computation. In *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE*, pages 418–419. IEEE, 2009. 23

[172] L. von Ahn, M. Blum, and J. Langford. Telling humans and computers apart automatically. *Communications of the ACM*, 47(2):56–60, 2004. 6

[173] L. von Ahn, B. Maurer, C. McMillen, D. Abraham, and M. Blum. RECAPTCHA: Human-based character recognition via web security measures. *Science*, 321(5895):1465, 2008. 6

[174] C. S. Wallace and D. M. Boulton. An information measure for classification. *Computer Journal*, 11(2):185–194, 1968. 20

[175] C. S. Wallace and D. L. Dowe. Minimum message length and Kolmogorov complexity. *Computer Journal*, 42(4):270–283, 1999. Special issue on Kolmogorov complexity. 20

[176] G. Wang, M. Mohanlal, C. Wilson, X. Wang, M. Metzger, H. Zheng, and B. Y. Zhao. Social Turing Tests: Crowdsourcing sybil detection. *arXiv preprint arXiv:1205.3856*, 2012. 6

[177] K. Warwick. Turing Test success marks milestone in computing history. *University or Reading Press Release*, 8 June 2014. 5

[178] C. J. C. H. Watkins and P. Dayan. Q-learning. *Mach. learning*, 8(3):279–292, 1992. 22

[179] D. J. Weiss. Better data from better measurements using computerized adaptive testing. *Journal of Methods and Measurement in the Social Sciences*, 2(1):1–27, 2011. 11

[180] J. Weizenbaum. ELIZA – a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):3645, 1966. 5

[181] M.P. Wellman, D.M. Reeves, K.M. Lochner, and Y. Vorobeychik. Price prediction in a trading agent competition. *J. Artif. Intell. Res. (JAIR)*, 21:19–36, 2004. 14

[182] D. R. White, J. McDermott, M. Castelli, L. Manzoni, B. W. Goldman, G. Kronberger, W. Jaśkowski, U.-M. O'Reilly, and S. Luke. Better GP benchmarks: Community survey results and proposals. *Genetic Programming and Evolvable Machines*, 14:3–29, 2013. 11

[183] S. Whiteson, B. Tanner, M. E. Taylor, and P. Stone. Protecting against evaluation overfitting in empirical reinforcement learning. In *Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), 2011 IEEE Symposium on*, pages 120–127. IEEE, 2011. 2, 4, 5, 7

[184] S. Whiteson, B. Tanner, and A. White. The Reinforcement Learning Competitions. *The AI magazine*, 31(2):81–94, 2010. 11

[185] P. L. Williams and R. D. Beer. Information dynamics of evolved agents. In *From Animals to Animats 11*, pages 38–49. Springer, 2010. 20, 23

[186] M. Winikoff and S. Cranefield. On the testability of bdi agent systems. *J. Artif. Intell. Res. (JAIR)*, 51:71–131, 2014. 4

[187] D. H. Wolpert. The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8(7):1341–1390, 1996. 4

[188] D. H. Wolpert. What the no free lunch theorems really mean; how to improve search algorithms. Technical report, Santa fe Institute Working Paper, 2012. 4

[189] D. H. Wolpert and W. G. Macready. No free lunch theorems for search. Technical report, Technical Report SFI-TR-95-02-010 (Santa Fe Institute), 1995. 4

[190] D. H. Wolpert and W. G. Macready. Coevolutionary free lunches. *Evolutionary Computation, IEEE Transactions on*, 9(6):721–735, 2005. 4

[191] R. Yonck. Toward a standard metric of machine intelligence. *World Future Review*, 4(2):61–70, 2012. 22

[192] Z. Zatuchna and A. Bagnall. Learning mazes with aliasing states: An LCS algorithm with associative perception. *Adaptive Behavior*, 17(1):28–57, 2009. 8